arXiv:1301.0722v1 [cs.CL] 4 Jan 2013

Good parts first - a new algorithm for approximate search in lexica and string databases

Stefan Gerdjikov, Stoyan Mihov, Petar Mitankin, Klaus U. Schulz

Abstract

We present a new efficient method for approximate search in electronic lexica. Given an input string (the pattern) and a similarity threshold, the algorithm retrieves all entries of the lexicon that are sufficiently similar to the pattern. Search is organized in subsearches that always start with an exact partial match where a substring of the input pattern is aligned with a substring of a lexicon word. Afterwards this partial match is extended stepwise to larger substrings. For aligning further parts of the pattern with corresponding parts of lexicon entries, more errors are tolerated at each subsequent step. For supporting this alignment order, which may start at any part of the pattern, the lexicon is represented as a structure that enables immediate access to any substring of a lexicon word and permits the extension of such substrings in both directions. Experimental evaluations of the approximate search procedure are given that show significant efficiency improvements compared to existing techniques. Since the technique can be used for large error bounds it offers interesting possibilities for approximate search in special collections of "long" strings, such as phrases, sentences, or book titles.

1 Introduction

The problem of approximate search in large lexica is central for many applications like spell checking, text and OCR correction [Kuk92, DHH+97], internet search [CB04, AK05, LH99], computational biology [Gus97] etc. In a common setup the problem may be formulated as follows: A large set of words/strings called the lexicon is given as a static background resource. Given an input string (the pattern), the task is to efficiently find all entries of the lexicon where the Levenshtein distance between pattern and entry does not exceed a fixed bound specified by the user. The Levenshtein distance [Lev66] is often replaced by related distances. In the literature, the problem has found considerable attention, e.g. [Ofl96, BYN98, BCP02, MS04].

Classical solutions to the problem [Ofl96] try to align the pattern P with suitable lexicon words in a strict left-to-right manner, starting at the left border of the pattern. The lexicon is represented as a trie or deterministic finite-state

automaton, which means that each prefix of a lexicon word is only represented once and corresponds to a unique path beginning at the start state. During the search, only prefixes of lexicon words are visited where the distance to a prefix P' of the pattern does not exceed the given bound b. As a filter mechanism that checks if these conditions are always met, Ukonnen's method [Ukk85] or Levenshtein automata [SM02] have been used. The main problem with this solution is the so-called "wall effect": if we tolerate b errors and start searching in the lexicon from left to right, then in the first b steps we have to consider *all* prefixes of lexicon words. Eventually, only a tiny fraction of these prefixes will lead to a useful lexicon word, which means that our exhaustive initial search represents a waste of time.

In order to avoid the wall effect, we need to find a way of searching in the lexicon such that during the initial alignment steps between pattern and lexicon words the number of possible errors is as small as possible. The ability to realize such a search is directly related to the way the lexicon is represented. In [MS04] we used two deterministic finite-state automata as a joint index structure for the lexicon. The first "forward" automaton represents all lexicon entries as before. The second "backward" automaton represents all reversed entries of the lexicon. Given an erroneous input pattern, we distinguished two subcases: (i) most of the discrepancies between the pattern and the lexicon word are in the first half of the strings; and (ii) most of the discrepancies are in the second half. We apply two subsearches. For subsearch (i) we use the forward automaton. During traversal of the first half of the pattern we tolerate at most $b/2$ errors. Then search proceeds by tolerating up to b errors. For subsearch (ii) the traversal is performed on the reversed automaton and the reversed pattern in a similar way – in the first half starting from the back only $b/2$ errors are allowed, afterwards the traversal to the beginning tolerates b errors. In [MS04] it was shown that the performance gain compared to the classical solution is enormous and at the same time no candidate is missed.

In this paper we present a method that can be considered as an extension of the latter. The new method uses ideas introduced in the context of approximate search in strings in [WM92, Mye94, BYN99, NBY99, NBY00]. Assume that the pattern can be aligned with a lexicon word with not more than b errors. Clearly, if we divide the pattern into $b + 1$ pieces, then at least one piece will *exactly* match the corresponding substring of a lexicon word in the answer set. In the new approach we first find the lexicon substrings that exactly match such a given piece of the pattern ("good parts first"). Afterwards we continue by extending this alignment, stepwise attaching new pieces on the left or right side. For the alignment of new pieces, more errors are tolerated at each step, which guarantees that eventually b errors can occur. Since at later steps the set of interesting substrings to be extended is already small the wall effect is avoided, it does not hurt that we need to tolerate more errors. For this kind of search strategy, a new representation of the lexicon is needed where we can start traversal at any point of a word. In our new approach, the lexicon is represented as symmetric compact directed acyclic word graph (SCDAWG) [BBH$^+$87, IHS$^+$01] - a bidirectional index structure where we (i) have direct access to every substring of a lexicon

word and (ii) can deterministically extend any such substring *both to the left and to the right* to larger substrings of lexicon words. This index structure can be seen as a part of a longer development of related index structures [BBH$^+$87, Sto95, Gus97, Bre98, Sto00, Maa00, IHS$^+$01, IHS$^+$05, MMW09] extending work on suffix tries, suffix trees, and directed acyclic word graphs (DAWGs) [Wei73, McC76, Ukk95, CS84, BBH$^+$85].

Our experimental results show that the new method is much faster than previous methods mentioned above. For small distance bounds it often comes close to the theoretical limit, which is defined as a (in practice merely hypothetical) method where precomputed solutions are used as output and no search is needed. In our evaluation we not only consider "usual" lexica with single-word entries. The method is especially promising for collections of strings where the typical length is larger than in the case of conventional single-word lexica. Note that given a pattern P and an error bound b, long strings in the lexicon have long parts that can be exactly aligned with parts of P. This explains why even for large error bounds efficient approximate search is possible. In our tests we used a large collection of book titles, and a list of 351,008 full sentences from MEDLINE abstracts as "dictionaries". In both cases, the speed up compared to previous methods is drastic. Future interesting application scenarios might include, e.g., approximate search in translation memories, address data, and related language databases.

The paper is structured as follows. We start with some formal preliminaries in Section 2. In Section 3 we present our method informally using an example. In Section 4 we give a formal description of the algorithm, assuming that an appropriate index structure for the lexicon with the above functionality is available. In Section 5 we describe the symmetric compact directed acyclic word graph (SCDAWG). Section 6 gives a detailed evaluation of the new method, comparing search times achieved with other methods. Experiments are based on various types of lexica, we also look at distinct variants of the Levenshtein distance. In the Conclusion we comment on possible applications of the new method in spelling correction and other fields. We also add remarks on the historical sources for the index structure used in this paper.

2 Technical Preliminaries

Words over a given finite alphabet Σ are denoted $P, U, V, W, \ldots$, symbols σ, σ_i denote letters of Σ. The *empty word* is written ε. If $W = \sigma_1 \cdots \sigma_n$, then W^{rev} denotes the *reversed word* $\sigma_n \cdots \sigma_1$. The i-th symbol σ_i of the word $W = \sigma_1 \cdots \sigma_n$ is denoted W_i. In what follows the terms *string* and *word* are used interchangeably. The *length* (number of symbols) of a word W is denoted $|W|$. We write $U \circ V$ or UV for the *concatenation* of the words $U, V \in \Sigma^*$. A string U is called a *prefix* (resp. *suffix*) of $W \in \Sigma^*$ iff W can be represented in the form $W = U \circ V$ (resp. $W = V \circ U$) for some $V \in \Sigma^*$. A string V is a *substring* of $W \in \Sigma^*$ iff W can be represented in the form $W = U_1 \circ V \circ U_2$ for some $U_1, U_2 \in \Sigma^*$. The set of all strings over Σ is denoted Σ^*, and the set

of the nonempty strings over Σ is denoted Σ^+. By a *lexicon* or *dictionary* we mean a finite nonempty collection $\mathcal{D}$ of words. The set of all substrings (resp. prefixes, suffixes) of words in $\mathcal{D}$ is denoted *Subs*($\mathcal{D}$) (resp. *Pref*($\mathcal{D}$), *Suf*($\mathcal{D}$)). The set of the reversed words from $\mathcal{D}$ is denoted $\mathcal{D}^{rev}$. The *size of the lexicon* $\mathcal{D}$ is $||\mathcal{D}|| := \sum_{W \in \mathcal{D}} |W|$.

Definition 2.1 A *deterministic finite-state automaton* is a quintuple

$$\mathcal{A} = (Q, \Sigma, s, \delta, F)$$

where Σ is a finite input alphabet, Q is a finite set of states, $s \in Q$ is the start state, $\delta : Q \times \Sigma \to Q$ is a partial transition function, and $F \subseteq Q$ is the set of final states.

If $\mathcal{A} = (Q, \Sigma, s, \delta, F)$ is a deterministic finite-state automaton, the extended partial transition function δ^* is defined as usual: for each $q \in Q$ we have $\delta^*(q, \varepsilon) = q$. For a string $W\sigma$ ($W \in \Sigma^*, \sigma \in \Sigma$) $\delta^*(q, W\sigma)$ is defined iff both $\delta^*(q, W) = p$ and $\delta(p, \sigma) = r$ are defined. In this case, $\delta^*(q, W\sigma) = r$. We consider the size of a deterministic finite-state automaton $\mathcal{A}$ to be linear in the number of states $|Q|$ plus the number of the transitions $|\{(p, \sigma, q) \mid \delta(p, \sigma) = q\}|$. Assuming that the size (number of symbols) of the alphabet Σ is treated as a constant, the *size* of $\mathcal{A}$ is $O(|Q|)$.

Definition 2.2 A *generalized deterministic finite-state automaton* is a quintuple $\mathcal{A} = (Q, \Sigma, s, \delta, F)$, where Q, Σ, s and F are as above and $\delta : Q \times \Sigma^+ \to Q$ is a partial function with the following property: for each $q \in Q$ and each $\sigma \in \Sigma$ there exists at most one $U \in \Sigma^*$ such that $\delta(q, \sigma U)$ is defined.

A transition $\delta(q, \sigma U) = p$ is called a *σ-transition* from q. The above condition then says that for each $q \in Q$ and each $\sigma \in \Sigma$ there exists at most one σ-transition. In what follows, σ-transitions of the above form are often denoted $q \stackrel{\sigma U}{\to} p$. Let

$$\mathcal{V} = \{(q, V, p) \mid q, p \in Q, \mathcal{A} \text{ has a transition } q \stackrel{V}{\to} p\}.$$

The size of the generalized deterministic finite-state automaton $\mathcal{A}$ is considered to be $O(|Q| + \sum_{(q,V,p) \in \mathcal{V}} |V|)$, which is not $O(|Q|)$ in general.

2.1 Suffix tries for lexica

The following definitions capture possible index structures for search in lexica. First, we define *the trie for a lexicon* $\mathcal{D}$ as a tree-shaped deterministic finite-state automaton. Each state of this automaton represents a unique prefix of lexicon words. The final states represent complete words. Second, the *suffix trie for* $\mathcal{D}$ is defined as the trie of all suffixes in $\mathcal{D}$.

Definition 2.3 Let $\mathcal{D}$ be a lexicon over the alphabet Σ. The *trie for* $\mathcal{D}$ is the deterministic finite-state automaton $Trie(\mathcal{D}) = (Q, \Sigma, q_\varepsilon, \delta, \{q_U \mid U \in \mathcal{D}\})$ where $Q = \{q_U \mid U \in Pref(\mathcal{D})\}$ is a set of states indexed with the prefixes in *Pref*($\mathcal{D}$) and $\delta(q_U, \sigma) = q_{U \circ \sigma}$ for all $U \circ \sigma \in Pref(\mathcal{D})$.

Obviously, the size of $Trie(\mathcal{D})$ is $O(||\mathcal{D}||)$. While tries support left-to-right search for words of the lexicon, the next index structure supports left-to-right search for substrings of lexicon words.

Definition 2.4 Let $\mathcal{D}$ as above. The *suffix trie for* $\mathcal{D}$ is the deterministic finite-state automaton $STrie(\mathcal{D}) := Trie(Suf(\mathcal{D}))$.

In general, the size of the suffix trie for $\mathcal{D}$ is $O(||Suf(\mathcal{D})||)$ and $||Suf(\mathcal{D})||$ is quadratic with respect to $||\mathcal{D}||$. For example, for every $n \in \mathbb{N}$ the number of states in $STrie(\{a^n b^n\})$ is $(n+1)^2$.

Bidirectional suffix tries. We now introduce a bidirectional index structure supporting both left-to-right search and right-to-left search for substrings of lexicon words. For $U \in \Sigma^*$ always q_U is a state in $STrie(\mathcal{D})$ iff $q_{U^{rev}}$ is a state in $STrie(\mathcal{D}^{rev})$. Hence, following Giegerich and Kurtz [GK97], from the two suffix tries $STrie(\mathcal{D})$ and $STrie(\mathcal{D}^{rev})$ we obtain one bidirectional index structure by identifying each pair of states $(q_U, q_{U^{rev}})$ from the two structures.

Definition 2.5 The *bidirectional suffix trie* for $\mathcal{D}$ is the tuple $BiSTrie(\mathcal{D}) := (Q, \Sigma, q_\varepsilon, \delta_L, \delta_R, F)$, where $(Q, \Sigma, q_\varepsilon, \delta_R, G') = STrie(\mathcal{D})$, $F := \{q_U \in Q \mid U \in \mathcal{D}\}$ and $\delta_L : Q \times \Sigma \to Q$ is the partial function such that $(Q^{rev}, \Sigma, q_\varepsilon, \delta_L^{rev}, G'') = STrie(\mathcal{D}^{rev})$ for $Q^{rev} = \{q_{U^{rev}} \mid q_U \in Q\}$ and $\delta_L^{rev}(q_{U^{rev}}, x) = \delta_L(q_U, x)$.

Example 2.6 The bidirectional suffix trie for $\mathcal{D} = \{\texttt{ear}, \texttt{lead}, \texttt{real}\}$, is shown in Figure 1.

As in the case of one-directional structures, the main problem is the size of the index. In general, the size of $BiSTrie(\mathcal{D})$ is quadratic in the size of the $\mathcal{D}$. The final structure, which will be presented in Section 5, can be considered as a compacted version of the bidirectional suffix trie.

Remark 2.7 It is known that a suffix tree for a lexicon $\mathcal{D}$ can be stored in space $O(||\mathcal{D}||)$ and built online in time $O(||D||)$, [Ukk95]. Suffix trees are compacted variants of suffix tries. In this paper we use compact directed acyclic word graphs [BBH$^+$87, IHS$^+$05] which are minimized variants of suffix trees.

2.2 Approximate search in lexica and Levenshtein filters

Definition 2.8 The *Levenshtein distance* between $V, W \in \Sigma^*$, denoted $d_L(V, W)$, is the minimal number of edit operations needed to transform V into W. *Edit operations* are the deletion of a symbol, the insertion of a symbol, and the substitution of a symbol by another symbol in Σ.

In what follows, $Id := \{\langle \sigma, \sigma \rangle \mid \sigma \in \Sigma\}$ is considered as a set of identity operations.

Definition 2.9 A *set of generalized weighted operations* is a pair (Op, w) where

1. $Op \subseteq \Sigma^* \times \Sigma^*$ is a finite set of operations such that $Id \subseteq Op$,

2. $w : Op \to \mathbb{N}$ assigns to each operation $op \in Op$ a nonnegative integer weight $w(op)$ such that $w(op) = 0$ iff $op \in Id$.

If $op = \langle X, Y \rangle$ represents an operation in Op, then $l(op)$, *the left side of the operation*, is defined as $l(op) = X$ and $r(op)$, *the right side of the operation*, is defined as $r(op) = Y$. The *width* of $op \in Op$ is $|l(op)|$.

Definition 2.10 Let (Op, w) be a set of generalized weighted operations. An *alignment* is an arbitrary sequence $\alpha = op_1 op_2 \dots op_n \in Op^*$ of operations $op_i \in Op$. The notions of left (right) side and weight are extended to alignments in a natural way:

$$l(\alpha) = l(op_1)l(op_2)\dots l(op_n)$$
$$r(\alpha) = r(op_1)r(op_2)\dots r(op_n)$$
$$w(\alpha) = \textstyle\sum_{i=1}^{n} w(op_i).$$

Note that Definition 2.10 does not permit overlapping of operations in the sequence. In our setting, operations that transform the left side into the right side are applied simultaneously. Formally, each sequence of operations representing an alignment is a string over the alphabet Op.

Definition 2.11 The generalized distance induced by a given set of generalized weighted operations (Op, w) is the function $d : \Sigma^* \times \Sigma^* \to \mathbb{N} \cup \{\infty\}$ which is defined as:

$$d(V, W) = \min\{w(\alpha) \mid \alpha \in Op^*, l(\alpha) = V \text{ and } r(\alpha) = W\}.$$

We say that $\alpha \in Op^*$ is an *optimal alignment of* V *and* W iff $l(\alpha) = V$, $r(\alpha) = W$ and $w(\alpha) = d(V, W)$.

Remark 2.12 In terms of Definition 2.11 we can represent the Levenshtein d_L as the distance induced by (Op_L, w_L) where $Op_L = (\Sigma \cup \varepsilon) \times (\Sigma \cup \varepsilon) \setminus \{\langle \varepsilon, \varepsilon \rangle\}$ and $w_L(op) = 1$ for all $op \notin Id$.

Remark 2.13 Given a set of generalized weighted operations (Op, w), dynamic programming can be used to efficiently compute $d(V, W)$ for strings V and W, [Ukk85, Ver88].

In this paper, we are interested in solutions for the following algorithmic problem (*"approximate search in lexica"*):

Let $\mathcal{D}$ be a fixed lexicon, let d denote a given generalized distance between words. For an input pattern $P \in \Sigma^*$ and a bound $b \in \mathbb{N}$, efficiently find all words $W \in \mathcal{D}$ such that $d(P, W) \leq b$.

Definition 2.14 Let $b \in \mathbb{N}$ denote a given bound. By a *Levenshtein filter* for bound b we mean any algorithm that takes as input two words $P, U \in \Sigma^*$ and decides

1. if there exists a string $V \in \Sigma^*$ such that $d_L(P, U \circ V) \leq b$,
2. if $d_L(P, U) \leq b$.

More generally, if d is any generalized distance, a *filter for* d for bound b is an algorithm that takes as input two words $P, U \in \Sigma^*$ and decides

1. if there exists a string $V \in \Sigma^*$ such that $d(P, U \circ V) \leq b$,
2. if $d(P, U) \leq b$.

Note that a filter for d for bound b does not depend on the lexicon $\mathcal{D}$.

The interest in filters of the above form relies on the observation that in approximate search in lexica we often face a given input pattern $P \in \Sigma^*$. When we traverse the lexicon, which is represented as a trie or automaton, we want to recognize at the earliest possible point if the current path, which represents a prefix U of a lexicon word, can *not* be completed to any word that is close enough to P (Decision Problem 1). When reaching a final state representing a word $W = U$ of the lexicon we want to check if W satisfies the bound (Decision Problem 2). In [Ofl96], the matrix based dynamic programming approach was used to realize a Levenshtein filter. In [SM02] we introduced the concept of a Levenshtein automaton, which represents a more efficient filter mechanism.

In what follows we make a more general use of filters. Our lexicon traversal below starts from a *substring* of a lexicon word, which is compared to a substring P of the pattern. In addition to steps where we extend substrings on the right using a filter of the above form, we also use steps where we extend substrings with new symbols on the left. In this situation we need to check for given $P, U \in \Sigma^*$ if there exists a string V such that $d(P, V \circ U) \leq b$. This means that with suitable extensions of U on the left we might reach an interesting alignment partner for P among the substrings of lexicon words.

Remark 2.15 Assume that we have an algorithm that, given a distance d induced by (Op, w) and a bound b, constructs a filter for extension steps on the right of the above form. We may build a second filter for the symmetric distance $d^{rev} := (\{op^{rev} \mid op \in Op\}, w^{rev})$ where $op^{rev} := (l(op)^{rev}, r(op)^{rev})$ and $w^{rev}(op^{rev}) := w(op)$ for $op \in Op$. Obviously, for given $P, U \in \Sigma^*$ there exists a string V such that $d(P, V \circ U) \leq b$ iff there exists a string V' such that $d^{rev}(P^{rev}, U^{rev} \circ V') \leq b$. Hence the second "reversed filter" can be used to control extension steps on the left.

The use of filters is directly related to the "wall effect". When the lexicon offers many possibilities for extending a given prefix or substring of a lexicon word, then the search space in a crucial way depends on the bound b of the filter that is used. When using a large bound, a large number of extensions has to be considered. Note that typically short prefixes/substrings have a very large number of extensions in the lexicon, while long prefixes/substrings often point to a unique entry. From this perspective, the problem discussed in the paper

can be rephrased: we are interested in a search strategy where the use of large bounds in filters is only necessary for large substrings at the end of the search. When we construct alignments between the pattern and lexicon words, we want to build "good parts" first.

3 Basic Idea

In this section we explain the idea of our algorithm using a small example. We also characterize the kind of resources needed to achieve its efficient implementation. Consider the dictionary

$$\mathcal{D} = \{\texttt{ear}, \texttt{real}, \texttt{lead}\}.$$

Suppose that for the pattern

$$P = \texttt{dread}$$

we want to find all words W in $\mathcal{D}$ such that $d_L(P, W) \leq 2$. The standard way to solve the problem is a left-to-right search in the lexicon, using a filter for bound 2. As described above, we want to avoid the use of a large filter bound at the beginning of the search. We next illustrate a first approach along these lines, which is then refined.

Let W in $\mathcal{D}$ such that $d_L(P, W) \leq 2$. When we split $P = \texttt{dread}$ into the three parts $\texttt{d}$, $\texttt{re}$, $\texttt{ad}$, then there must be a corresponding representation of W in the form $W = W_1 \circ W_2 \circ W_3$ such that $d_L(\texttt{d}, W_1) + d_L(\texttt{re}, W_2) + d_L(\texttt{ad}, W_3) \leq 2$. We distinguish three cases, $d_L(\texttt{d}, W_1) = 0$, $d_L(\texttt{re}, W_2) = 0$, or $d_L(\texttt{ad}, W_3) = 0$. This leads to the following three subtasks:

1. Check if $\texttt{d}$ represents a substring of a word in $\mathcal{D}$. In the positive case, look for extensions V of $\texttt{d}$ on the right to words of the form $\texttt{d} \circ V \in \mathcal{D}$ such that $d_L(\texttt{dread}, \texttt{d}V) \leq 2$.

2. Check if $\texttt{re}$ represents a substring of a word in $\mathcal{D}$. In the positive case, look for extensions V_2 of $\texttt{re}$ on the right and extensions V_1 of $\texttt{re}V_2$ on the left to words of the form $V_1 \circ \texttt{re} \circ V_2 \in \mathcal{D}$ such that $d_L(\texttt{dread}, V_1\texttt{re}V_2) \leq 2$.

3. Check if $\texttt{ad}$ represents a substring of a word in $\mathcal{D}$. In the positive case, look for extensions V of $\texttt{ad}$ on the left to words of the form $V \circ \texttt{ad} \in \mathcal{D}$ such that $d_L(\texttt{dread}, V\texttt{ad}) \leq 2$.

The above task can be solved using an appropriate bidirectional index structure. As an illustration[1] we use the bidirectional suffix trie (cf. Def. 2.5) for $\mathcal{D} = \{\texttt{ear}, \texttt{lead}, \texttt{real}\}$, which is shown in Figure 1. The nodes q_U of the graph depicted correspond to the substrings U of our lexicon $\mathcal{D}$, nodes marked with

[1] We should stress that $BiSTrie(\mathcal{D})$ is just used for illustration purposes. In general, the size of $BiSTrie(\mathcal{D})$ is quadratic in the size of the $\mathcal{D}$, which means that a more condensed structure is needed in practice.

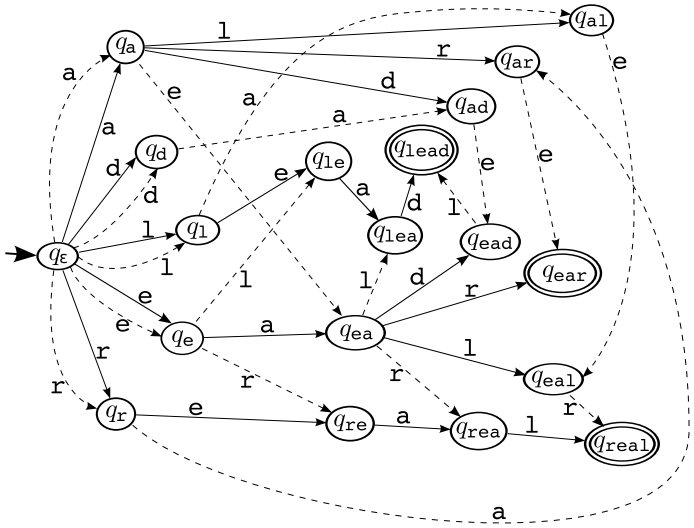


Figure 1: The bidirectional suffix trie $BiSTrie(\mathcal{D})$ for $\mathcal{D} = \{\mathtt{ear}, \mathtt{lead}, \mathtt{real}\}$ represents all substrings of $\mathcal{D}$ and allows extending each substring either to the right by following the *solid arcs* or to the left by following the *dashed arcs*.

a double ellipse represent words in $\mathcal{D}$. Following the *solid arcs* we extend the current substring to the right. Starting from q_ε and traversing solid arcs we find any substring. If we follow the *dashed arcs* we extend the current substring to the left.

It should be obvious how we may use the graph to solve the three subtasks in our example mentioned above. As an example, we consider Subtask 2. Using the index we see that `re` is a substring of a word in $\mathcal{D}$. Right extension steps of `re` in the index are controlled using a Levenshtein filter for pattern suffix `read` and bound 2. We find the two extensions `rea` and `real`. Then, for the left extension steps we use the filter for the full pattern `dread` and bound 2. The index shows that both `rea` and `real` cannot be extended on the left. However, since already $d_L(\mathtt{dread}, \mathtt{rea}) \leq 2$ and $d_L(\mathtt{dread}, \mathtt{real}) \leq 2$ the filter licenses the empty left extension. Among the two resulting substrings, $\mathtt{real} \in \mathcal{D}$ is a solution. In a similar way, solving Subtask 3 leads to the second solution `lead`. When we abstract from our small example, the above procedure gives rise to the following

First search idea. Split $P = P_1 \circ \cdots \circ P_{b+1}$ into $b+1$ parts P_i of approximately the same length and apply $b+1$ subsearches. For the i-th subsearch, first check if P_i is a substring of a lexicon word (Step 1). In the positive case, try to extend P_i to larger substrings of lexicon words, using a Levenshtein filter for bound b (Step 2).

A nice aspect of this a search strategy is that each subsearch starts with an exact match (Step 1), which represents a search with filter bound 0. However,

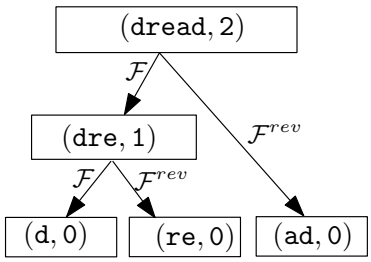


Figure 2: Reducing the original query (dread, 2) into simpler ones. As a result we obtain an ordered binary tree representing search alternatives. The labels of the arcs indicate what sort of filter has to be used at the extension steps. The label $\mathcal{F}$ shows that we extend to the right and thus an ordinary filter is required, whereas the label $\mathcal{F}^{rev}$ means that we extend to the left and therefore a reverse filter (see Remark 2.15) has to supervise this step. The bound that determines a filter coincides with the threshold of the query written in the parent node.

afterwards in Step 2 we immediately use a Levenshtein filter for the full bound b for all left and right extension steps. If b is large, this may lead to a large search space.

Improved search idea. We now look for a refinement where we can use small filter bounds for the initial extension steps. To this end, we first slightly generalize the problem and search all *substrings* V of words in $\mathcal{D}$ such that $d_L(P, V) \leq b$. Afterwards we simply filter those substrings that represent entries of $\mathcal{D}$.

We illustrate the improved search procedure using again our small example. In what follows, the notation (dread, 2) is used as a shorthand for the algorithmic task to find all substrings $V \in \mathit{Subs}(\mathcal{D})$ such that $d_L(\texttt{dread}, V) \leq 2$, and similarly for other strings and bounds. The expression (dread, 2) is called a *query* with query pattern dread and bound 2. Now consider the query tree depicted in Figure 2. The idea is to solve the problems labeling the nodes in a bottom-up manner. The three leaves exactly correspond to the Steps 1 in the three subtasks discussed above: in fact, to solve the problems (d, 0), (re, 0) and (ad, 0) just means to check if d, re, or ad are substrings of lexicon words. We then solve problem (dre, 1). This involves two independent steps.

1. We look for extensions of the substring d (as a solution of the left child in the tree) at the right.

2. We look for extensions of the substring re (as a solution of the right child in the tree) at the left.

It is important to note that both extension steps are controlled using a Levenshtein filter for bound 1 for $P' =$ dre (see Figure 2). As a result we obtain the single solution re for the query (dre, 1). The next step in the bottom-up

procedure looks at the root node (`dread`, 2). Solving this node again involves two independent steps.

1. We look for extensions of the substring `re` (as a solution of the left child in the tree) at the right.
2. We look for extensions of the substring `ad` (as a solution of the right child in the tree) at the left.

At this final step we cannot avoid the use of a Levenshtein filter for `dread` and bound 2. We respectively obtain (1) `rea`, `real` and (2) `dead`, `lead`.

Comparing the two search strategies, we see that at least Subtasks 1 and 2 have been replaced by a subsearch where we use filter bound 2 only at the last extension step where we already found `re` and want to solve (`dread`, 2) adding right extensions. More generally, search trees of this form offer a possibility to postpone the use of large filter bounds to the end of the search. Details will be given in the next section where we formally describe the refined procedure.

Remark 3.1 In order to efficiently realize a bottom-up subsearch of the form indicated above we need

1. an index structure that supports the following tasks:
 (a) given a string V, efficiently decide if V represents a substring of a lexicon word,
 (b) given a substring V of a lexicon word, give immediate access to all substrings of lexicon words of the form $V \circ \sigma$ that add one letter $\sigma \in \Sigma$ to the right,
 (c) given a substring V of a lexicon word, give immediate access to all substrings of lexicon words of the form $\sigma \circ V$ that add one letter $\sigma \in \Sigma$ to the left.
2. A filter for the bound b specified at the parent node faced at an upward step. The filter takes as first input the query pattern P' specified at the parent node. Subsearches start with a given solution of the left (right) child query. When adding letters to the right (left) we use a conventional ("reversed") filter, cf. Remark 2.15.

Remark 3.2 A similar idea was introduced by Navaro and Baeza-Yates, [NBY00], for approximate search of a pattern in the set of substrings of a long text. In [NBY00] the authors use suffix arrays for its realization and analyze how to organize the splitting to optimize the efficiency of this approach in terms of the length of the text and pattern. Their theoretical results show that this technique improves over the naive algorithm in some cases, but still it does not avoid the wall effect in general. In [NBY99] the same authors present an algorithm for online approximate search of substrings of a long text. Their algorithm, as the algorithm presented here, uses binary trees representing search alternatives to reduce the search space. The essential difference is that their algorithm is online, i.e. does not rely on a precomputed index.

4 Search procedure

The purpose of this section is to provide a formal description of the approach considered in the previous section. In what follows we assume that $\mathcal{D}$ is a fixed lexicon and d is a given generalized distance induced by (Op, w). For input strings $P \in \Sigma^*$ and a bound $b \in \mathbb{N}$ we want to retrieve all words $W \in \mathcal{D}$ such that $d(P, W) \leq b$. We consider the case where each operation $op \in Op$ has width ≤ 1 (cf. Def. 2.9). In the Appendix we show how essentially the same technique can be used for arbitrary generalized distances.

Definition 4.1 A *query* is a pair (P', b') where $P' \in \Sigma^*$ is a substring of P and $b' \leq b$. The set $Sol_D(P', b') := \{V \in Subs(\mathcal{D}) \mid d(P', V) \leq b'\}$ is called the *solution set* for (P', b').

The search procedure has three phases. We first build a search tree for the query (P, b). Then, using a bottom-up procedure we solve all queries of the search tree, in particular (P, b). The final step is trivial. We simply select from $Sol_D(P, b)$ those elements that represent entries of $\mathcal{D}$.

4.1 Building the search tree for a pattern

We explain how to obtain for a given pattern P a binary tree $\mathcal{T}_P$ with queries assigned to each node, see Figure 2.

Select any rooted ordered tree $\mathcal{T}$ with $b+1$ leaves $\lambda_1, \ldots, \lambda_{b+1}$ (enumerated in canonical left-to-right ordering) where each non-leaf node has exactly 2 children. Then decorate the nodes of $\mathcal{T}$ with queries to define the search tree $\mathcal{T}_P$: Split the pattern P in the form $P = P_1 \circ P_2 \cdots \circ P_{b+1}$ where the P_i are substrings of P of almost equal length, i.e. $||P_i| - |P_j|| \leq 1$ $(1 \leq i, j \leq b+1)$. To each leaf λ_i of $\mathcal{T}$ assign the query $(P_i, 0)$. To each non-leaf node η of $\mathcal{T}$ assign the query $(P_i \circ \cdots \circ P_{i+b'}, b')$ where $\lambda_i, \ldots, \lambda_{i+b'}$ is the sequence of leaves representing descendants of η in $\mathcal{T}$ in the natural left-to-right ordering. (Note that the root of $\mathcal{T}_P$ has label (P, b), which is the original query.)

Example 4.2 In the example considered in Section 3 we had $b = 2$, $P = \texttt{dread}$, $P_1 = \texttt{d}$, $P_2 = \texttt{re}$, $P_3 = \texttt{ad}$. As our starting point $\mathcal{T}$ for decoration, we selected one among two possible binary rooted ordered trees.

Remark 4.3 The choice of a tree $\mathcal{T}$ satisfying the above conditions influences the time needed to solve the query. The general philosophy is to avoid queries of the form (P', b') where P' is a short word and b' is a large bound. A good choice is the use of a balanced tree where all paths from the root reach a certain length. Other optimizations represent a possible subject for future studies.

4.2 Computation of solution sets

For each query (P', b') of the tree $\mathcal{T}_P$ we compute a set $S_D(P', b')$ in a bottom-up fashion. We shall prove below that $S_D(P', b')$ is the solution set $Sol_D(P', b')$ in each case.

Initialization steps. For a leaf query $(P_i, 0)$ we decide if P_i is a substring of a lexicon word. In the positive case we let $S_D(P_i, 0) := \{P_i\}$, otherwise we define $S_D(P_i, 0) := \emptyset$.

Definition 4.4 *Extension steps.* Let (P', b') denote the query at a non-leaf node η of $\mathcal{T}_P$, let (P'_1, b'_1) and (P'_2, b'_2) denote the queries of the two children η_1, η_2 of η, which are given in the natural left-to-right ordering. Given the sets $S_D(P'_1, b'_1)$ and $S_D(P'_2, b'_2)$ we define $S_D(P', b')$ as the union of the two sets S_1 and S_2 defined as

$$\begin{aligned} S_1 &:= \{U \circ V \in \mathit{Subs}(\mathcal{D}) \mid U \in S_D(P'_1, b'_1), d(P'_1, U) + d(P'_2, V) \leq b'\} \\ S_2 &:= \{V \circ U \in \mathit{Subs}(\mathcal{D}) \mid U \in S_D(P'_2, b'_2), d(P'_1, V) + d(P'_2, U) \leq b'\} \end{aligned}$$

Proposition 4.5 *The computation of solution sets is correct: for each query (P', b') of $\mathcal{T}_P$ we have $S_D(P', b') = \mathit{Sol}_D(P', b')$.*

4.3 Correctness proof and remarks

To prove Proposition 4.5, some preparations are needed.

Remark 4.6 Let η denote a non-leaf node of $\mathcal{T}_P$ decorated with query (P', b'). Let (P'_1, b'_1) and (P'_2, b'_2) denote the queries of the two children η_1, η_2 of η, which are given in the natural left-to-right ordering. Then we have $P' = P'_1 \circ P'_2$ and $b'_1 + b'_2 = b' - 1$.

Proposition 4.7 *Let $\mathcal{D}$ and $d = (\mathrm{Op}, w)$ as above, assume that each operation in Op has width ≤ 1, let $b \in \mathbb{N}$. If $P \in \Sigma^*$ is a word with $P = P_1 \circ P_2$ and $\alpha \in \mathrm{Op}^*$ is an alignment with $l(\alpha) = P$ and weight $w(\alpha) \leq b$, then*

1. *α can be represented in the form $\alpha = \alpha_1 \circ \alpha_2$ such that $l(\alpha_1) = P_1$ and $l(\alpha_2) = P_2$.*
2. *if b' and b'' are two arbitrary integers with the property $b' + b'' = b - 1$, then $w(\alpha_1) \leq b'$ or $w(\alpha_2) \leq b''$.*

Proof. Since α is a sequence of operations $\alpha = op_1 \dots op_n$ and $P = l(op_1) \dots l(op_n)$, the first part follows immediately from the fact that $|l(op_i)| \leq 1$ $(1 \leq i \leq n)$. The second statement is an obvious consequence. □

Corollary 4.8 *If P and W are arbitrary words with $d(P, W) \leq b$ and $P = P_1 \circ P_2$, then:*

1. *W can be represented in the form $W = W_1 \circ W_2$ such that $d(P, W) = d(P_1, W_1) + d(P_2, W_2)$.*
2. *if $b' + b'' = b - 1$, then $d(P_1, W_1) \leq b'$ or $d(P_2, W_2) \leq b''$,*

Proof. Let α be an optimal alignment of P and W. Then $w(\alpha) = d(P, W) \leq b$. We can define $W_i = r(\alpha_i)$ for $i = 1, 2$ where α_i are the alignments provided by Proposition 4.7. The second statement follows since $w(\alpha_i) \geq d(P_i, W_i)$, by the definition of a d-distance. □

(Proof of Proposition 4.5.) This is obvious for the leaf queries. Consider a non-leaf node η of $\mathcal{T}_P$ with query (P', b'), let (P'_1, b'_1) and (P'_2, b'_2) denote the queries of the two children η_1, η_2 of η, which are given in the natural left-to-right ordering. We may assume that $S_D(P'_i, b'_i) = Sol_D(P'_i, b'_i)$ for $i = 1, 2$. Remark 4.6 shows that $P' = P'_1 \circ P'_2$. Consider an element $U \circ V$ of S_1 (we use the notation introduced in Section 4.2). We have $d(P', U \circ V) = d(P'_1 \circ P'_2, U \circ V) \leq d(P'_1, U) + d(P'_2, V) \leq b'$ which shows that $U \circ V \in Sol_D(P', b')$. Hence $S_1 \subseteq Sol_D(P', b')$. Similarly we see that $S_2 \subseteq Sol_D(P', b')$. Conversely consider an element $W \in Sol_D(P', b')$. Since $P' = P'_1 \circ P'_2$, Corollary 4.8 shows that W can be represented as $W = W_1 \circ W_2$ such that $d(P'_1, W_1) + d(P'_2, W_2) = d(P', W)$ and we have (1) $d(P'_1, W_1) \leq b'_1$ or (2) $d(P'_2, W_2) \leq b'_2$. In case (i), $W_i \in S_D(P'_i, b'_i) = Sol_D(P'_i, b'_i)$, which shows that $W = W_1 \circ W_2$ is found in S_i $(i = 1, 2)$. Hence $Sol_D(P', b') \subseteq S_1 \cup S_2 = S_D(P', b')$. □

Remark 4.9 It is simple to see that in the example presented in the previous section the computation of solutions sets follows exactly the above procedure. Following the definition of the Extension steps, we need to construct the complete sets of candidates $S_D(P'_1, b'_1)$ and $S_D(P'_2, b'_2)$ in order to compute the sets S_1 and S_2 for the query node (P', b') and eventually determine $S_D(P', b') = S_1 \cup S_2$. This corresponds to a bottom-up traversal of the search tree $\mathcal{T}_P$.

Remark 4.10 Observe that given a candidate $U_1 \in S_D(P'_1, b'_1)$, the set of successful candidates $U_1 \circ V_1 \in S_1$ with $d(U_1, P'_1) + d(V_1, P'_2) \leq b'$ which result in the right extension steps depend only on the specific candidate U_1, the dictionary $\mathcal{D}$ and the query node (P', b') but not on $S_D(P'_1, b'_1)$. A similar observation is valid for the successful candidates $U_2 \in S_D(P'_2, b'_2)$.

Remark 4.11 Remark 4.10 means that the target set $S(P, b)$ can be constructed by using any traversal algorithm $\mathcal{A}$ of the search tree $\mathcal{T}(P)$ which satisfies the following three conditions:

1. it correctly initializes the candidate sets $S_D(P_i, 0)$ for each leaf $(P_i, 0)$.
2. if $\mathcal{A}$ generates a candidate U_1 in a node (P'_1, b'_1) which is a left child of (P', b'), then $\mathcal{A}$ generates also all successful candidates $U_1 \circ V_1$ for the node (P', b') such that $d(U_1, P'_1) + d(V_1, P'_2) \leq b'$.
3. if $\mathcal{A}$ generates a candidate U_2 in a node (P'_2, b'_2) which is a right child of (P', b'), then $\mathcal{A}$ generates also all successful candidates $U_1 \circ V_1$ for the node (P', b') such that $d(U_1, P'_1) + d(V_1, P'_2) \leq b'$.

In particular one can replace the bottom-up traversal by a depth first search algorithm.

Remark 4.12 It is obvious to see that the efficient realization of the above search algorithm can be based on the resources described in Remark 3.1. The efficient computation of the sets S_1 and S_2 in the bottom-up steps is achieved by using the given index structure for extensions on the right and left, respectively. Each extension by a single letter is controlled using a filter for the generalized distance d for the appropriate bound. As we mentioned earlier, the index structure shown in Figure 1 only serves for illustration purposes. When using this construction there is a one-to-one correspondence between the nodes and the substrings of lexicon words. In general, the number of substrings of entries in $\mathcal{D}$ is quadratic in the size (number of symbols) of $\mathcal{D}$. In the next section we shall describe an index structure that has the same functionality and needs storage space linear in the size (number of symbols) of the lexicon $\mathcal{D}$.

Remark 4.13 The approach that we proposed in this section is closely related to the algorithm of Myers [Mye94] for approximate search in strings. The main difference is that we have a fixed threshold b for the number of errors, whereas in [Mye94] the threshold is given in terms of percentage of symbols. This imposes different ways of handling the arising situation and modifications related with the application of the pigeonhole principle. Thus for query words of length at least $b + 1$ we shall always have an initialization with an exact match, whereas in Myers' situation this assumption is not obligatory fulfilled and he is not able to use it.

5 Symmetric compact directed acyclic word graphs

In this section we describe the bidirectional index structure for search in the lexicon. Afterwards we explain how the index structure supports the computation of solutions sets described in Section 4.2 during online steps. As before, $\mathcal{D}$ denotes the given lexicon, $\#\mathcal{D}\$$ denotes the variant where the new symbols $\#$ and $\$$ are attached as the first and the last symbol to each lexicon word, $\Sigma_{\#\$} = \Sigma \cup \{\#, \$\}$.

5.1 The index structure

In Section 2.1 we described how suffix tries and suffix tries for reversed words of the lexicon can be merged into a bidirectional (quadratic) index structure, using a bijective correspondence between the states of the two substructures. We now introduce a bidirectional index of linear size, which is used in our method for approximate search. Furthermore we present a new algorithm for online construction of such index. To build the index we crucially use an algorithm from [IHS+05] for online construction of one-directional compact directed acyclic word graphs.

Definition 5.1 Let $W \in \Sigma^*_{\#\$}$, let $0 \leq i \leq |W|$. A word $V \in \Sigma^*_{\#\$}$ of length $0 < |V| \leq |W|$ is said to *start at position i in W* if $i + |V| - 1 \leq |W|$ and $W_i W_{i+1} \ldots W_{i+|V|-1} = V$. Similarly V is said to *end at position i in W* if

$0 < i - |V| + 1$ and $W_{i-|V|+1}W_{i-|V|+2}\ldots W_i = V\}$. We define the functions $startpos_W$ and $endpos_W$ as

$$\begin{aligned} startpos_W(V) &:= \{i \in \mathbb{N} \mid V \text{ starts in } W \text{ at position } i\}, \\ endpos_W(V) &:= \{i \in \mathbb{N} \mid V \text{ ends in } W \text{ at position } i\}. \end{aligned}$$

In addition, let $startpos_W(\varepsilon) := endpos_W(\varepsilon) := \{0, 1, \ldots, |W|\}$.
We define the equivalence relations $\overrightarrow{\equiv}_W$ and $\overleftarrow{\equiv}_W$ on $\Sigma^*_{\#\$}$ as:

$$\begin{aligned} X \overrightarrow{\equiv}_W Y &\iff startpos_W(X) = startpos_W(Y), \\ X \overleftarrow{\equiv}_W Y &\iff endpos_W(X) = endpos_W(Y). \end{aligned}$$

Definition 5.2 The equivalence relations $\overrightarrow{\equiv}_{\#\mathcal{D}\$}$ and $\overleftarrow{\equiv}_{\#\mathcal{D}\$}$ are defined on $Subs(\#\mathcal{D}\$)$ as follows. For every $X, Y \in Subs(\#\mathcal{D}\$)$

$$X \overrightarrow{\equiv}_{\#\mathcal{D}\$} Y \iff \forall W \in \#\mathcal{D}\$: X \overrightarrow{\equiv}_W Y,$$

$$X \overleftarrow{\equiv}_{\#\mathcal{D}\$} Y \iff \forall W \in \#\mathcal{D}\$: X \overleftarrow{\equiv}_W Y.$$

In what follows, the equivalence class of a substring $V \in Subs(\#\mathcal{D}\$)$ w.r.t. $\overrightarrow{\equiv}_{\#\mathcal{D}\$}$ ($\overleftarrow{\equiv}_{\#\mathcal{D}\$}$) is written $\overrightarrow{[V]}$ ($\overleftarrow{[V]}$). It is easy to prove the following properties of the function $startpos_W$ ($endpos_W$).

Proposition 5.3 *Let $W, X, Y \in \Sigma^*_{\$}$ be arbitrary strings.*

1. *If $startpos_W(X) \cap startpos_W(Y) \neq \emptyset$ ($endpos_W(X) \cap endpos_W(Y) \neq \emptyset$), then X is a prefix (suffix) of Y or vice versa,*
2. *If Y is a prefix (suffix) of X, then $startpos_W(X) \subseteq startpos_W(Y)$ ($endpos_W(X) \subseteq endpos_W(Y)$).*

Proposition 5.3 can be used to show that for any two elements $X, Y \in \overrightarrow{[V]}$ ($X, Y \in \overleftarrow{[V]}$) either X is a prefix (suffix) of Y or vice versa. Consequently we can define *the canonical representative* $\overrightarrow{X}$ of $\overrightarrow{[X]}$ ($\overleftarrow{X}$ of $\overleftarrow{[X]}$) as the *longest word* $\overrightarrow{X} \in \overrightarrow{[X]}$ ($\overleftarrow{X} \in \overleftarrow{[X]}$).

Proposition 5.4 *[IHS+05] For every $X \in Subs(\#\mathcal{D}\$)$ there uniquely exist $\alpha, \beta \in \Sigma^*_{\#\$}$ such that $\overleftarrow{X} = \alpha X$ and $\overrightarrow{X} = X\beta$.*

Definition 5.5 For every $X \in Subs(\#\mathcal{D}\$)$ we define $\overleftrightarrow{X} := \alpha X \beta$ where $\overleftarrow{X} = \alpha X$ and $\overrightarrow{X} = X\beta$. The equivalence relation $\overleftrightarrow{\equiv}_{\#\mathcal{D}\$}$ is defined on $Subs(\#\mathcal{D}\$)$ as follows. For every $X, Y \in Subs(\#\mathcal{D}\$)$

$$X \overleftrightarrow{\equiv}_{\#\mathcal{D}\$} Y \iff \overleftrightarrow{X} = \overleftrightarrow{Y}.$$

In what follows the equivalence class of X w.r.t. $\overleftrightarrow{\equiv}_{\#\mathcal{D}\$}$ is written $\overleftrightarrow{[X]}$. We shall use $\overleftarrow{\equiv}$, $\overrightarrow{\equiv}$ and $\overleftrightarrow{\equiv}$ as shorthands for $\overleftarrow{\equiv}_{\#\mathcal{D}\$}$, $\overrightarrow{\equiv}_{\#\mathcal{D}\$}$ and $\overleftrightarrow{\equiv}_{\#\mathcal{D}\$}$ respectively.

Proposition 5.6 *[BBH+87] The equivalence relation $\overleftrightarrow{\equiv}$ is the transitive closure of $\overrightarrow{\equiv}$ and $\overleftarrow{\equiv}$.*

Proposition 5.7 *The equivalence relation $\overleftarrow{\equiv}$ is right-invariant: for arbitrary substrings $X, Y \in Subs(\#\mathcal{D}\$)$ and arbitrary extensions of the form $X \circ U, Y \circ U \in Subs(\#\mathcal{D}\$)$ always $X \overleftarrow{\equiv} Y$ implies $X \circ U \overleftarrow{\equiv} Y \circ U$. The equivalence relation $\overrightarrow{\equiv}$ is left-invariant.*

Definition 5.8 [BBH+87] The *directed acyclic word graph* (DAWG) for $\#\mathcal{D}\$$ is the deterministic finite-state automaton $\overleftarrow{\mathcal{A}}(\#\mathcal{D}\$) := (Q_{\overleftarrow{\mathcal{A}}}, \Sigma_{\#\$}, \overleftarrow{[\varepsilon]}, \delta_{\overleftarrow{\mathcal{A}}}, F_{\overleftarrow{\mathcal{A}}})$ where

- $Q_{\overleftarrow{\mathcal{A}}}$ is the set of all equivalence classes $\overleftarrow{[V]}$ w.r.t. $\overleftarrow{\equiv}$,
- the start state is $\overleftarrow{[\varepsilon]} \in Q_{\overleftarrow{\mathcal{A}}}$,
- the (partial) transition function $\delta_{\overleftarrow{\mathcal{A}}}$ is defined as $\delta_{\overleftarrow{\mathcal{A}}}(\overleftarrow{[V]}, \sigma) = \overleftarrow{[V \circ \sigma]}$ for all substrings $V \circ \sigma$ of $\#\mathcal{D}\$$,
- the set of final states is $F_{\overleftarrow{\mathcal{A}}} := \{\overleftarrow{[V]} \mid \overleftarrow{[V]} \cap \#\mathcal{D}\$ \neq \emptyset\}$.

Note that the right-invariance of $\overleftarrow{\equiv}$ implies that $\overleftarrow{\mathcal{A}}$ is well-defined. The DAWG for $\#\mathcal{D}\$$ can be used (i) to check if a string V is in $Subs(\#\mathcal{D}\$)$ and in the positive case (ii) to check in constant time if a right extension $V \circ \sigma$ again represents such a substring: for solving problem (i) we start a traversal of $\overleftarrow{\mathcal{A}}(\#\mathcal{D}\$)$ from $\overleftarrow{[\varepsilon]}$ with the letters of V. Then $V \in Subs(\#\mathcal{D}\$)$ iff all transitions are defined. In the positive case the traversal leads to the state $\overleftarrow{[V]}$. To solve problem (ii) we check if $\delta_{\overleftarrow{\mathcal{A}}}(\overleftarrow{[V]}, \sigma)$ is defined. Note that for tasks (i) and (ii) we need not fix a set of final states. With the above definition of final states we may check if a substring of the form $\#V\$$ represents a full entry of the lexicon. This holds iff $\delta^*_{\overleftarrow{\mathcal{A}}}(\overleftarrow{[\varepsilon]}, \#V\$)$ is (defined and) final. Analogously, using $\overrightarrow{\equiv}$, we define $\overrightarrow{\mathcal{A}}(\#\mathcal{D}\$)$, which can be used to check for left extensions $\sigma \circ V$. The question is how to merge $\overleftarrow{\mathcal{A}}$ and $\overrightarrow{\mathcal{A}}$ into one bidirectional index.

Definition 5.9 [BBH+87] The *compact directed acyclic word graph* (CDAWG) for $\#\mathcal{D}\$$ is the generalized deterministic finite-state automaton $\overleftarrow{\mathcal{C}}(\#\mathcal{D}\$) := (Q_{\overleftrightarrow{\mathcal{C}}}, \Sigma_{\#\$}, \overleftrightarrow{[\varepsilon]}, \delta_{\overleftrightarrow{\mathcal{C}}}, F_{\overleftrightarrow{\mathcal{C}}})$ where

- $Q_{\overleftrightarrow{\mathcal{C}}}$ is the set of all equivalence classes $\overleftrightarrow{[V]}$ w.r.t. $\overleftrightarrow{\equiv}$,
- the start state is $\overleftrightarrow{[\varepsilon]} \in Q_{\overleftrightarrow{\mathcal{C}}}$,
- the (partial) transition function $\delta_{\overleftrightarrow{\mathcal{C}}} : Q_{\overleftrightarrow{\mathcal{C}}} \times \Sigma^+_{\#\$} \to Q_{\overleftrightarrow{\mathcal{C}}}$ is defined for a string of the form σU ($\sigma \in \Sigma_{\#\$}$, $U \in \Sigma^*_{\#\$}$) iff $\overleftrightarrow{V \circ \sigma} = \overleftrightarrow{V} \circ \sigma U$. The value is
$$\delta_{\overleftrightarrow{\mathcal{C}}}(\overleftrightarrow{[V]}, \sigma U) = \overleftrightarrow{[V \circ \sigma]}$$

- the set of final states is $F_{\overleftrightarrow{\mathcal{C}}} := \{\overleftrightarrow{[V]} \mid \overleftrightarrow{[V]} \cap \#\mathcal{D}\$ \neq \emptyset\}$.

$\overleftarrow{\mathcal{C}}(\#\mathcal{D}\$)$ can be considered as a compacted variant of $\overleftarrow{\mathcal{A}}(\#\mathcal{D}\$)$, because $\overleftarrow{\mathcal{C}}(\#\mathcal{D}\$)$ can be obtained from $\overleftarrow{\mathcal{A}}(\#\mathcal{D}\$)$ by replacing chains of the type $q_0 \stackrel{\sigma_1}{\to} q_1 \stackrel{\sigma_2}{\to} q_2 \ldots \stackrel{\sigma_n}{\to} q_n$ with multi-letter transitions $q_0 \stackrel{\sigma_1\sigma_2\ldots q_n-1\sigma_n}{\longrightarrow} q_n$ iff states q_i for $1 < i < n$ are implicit[2] and q_1 and q_n are explicit, [BBH+87]. Analogously we define $\overrightarrow{\mathcal{C}}(\#\mathcal{D}\$)$ with (partial) transition function $\delta_{\overrightarrow{\mathcal{C}}}$:

$$\delta_{\overrightarrow{\mathcal{C}}}(\overleftrightarrow{[V]}, \sigma U) = \overleftarrow{[\sigma \circ \overleftrightarrow{V}]} \iff \overleftarrow{\sigma \circ \overleftrightarrow{V}} = U^{rev}\sigma \circ \overleftrightarrow{V}.$$

$\overrightarrow{\mathcal{C}}(\#\mathcal{D}\$)$ can be considered as well as a compacted variant of $\overrightarrow{\mathcal{A}}(\#\mathcal{D}\$)$ in the above sense. Note that both automata have the same set of states. Hence $\overleftarrow{\mathcal{C}}(\#\mathcal{D}\$)$ and $\overrightarrow{\mathcal{C}}(\#\mathcal{D}\$)$ are naturally merged into one bidirectional index. The following index is used in our method for approximate search.

Definition 5.10 [IHS+01] The bidirectional *symmetric compact acyclic word graph* (SCDAWG) for $\#\mathcal{D}\$$ is

$$\overleftrightarrow{\mathcal{C}}(\#\mathcal{D}\$) := (Q_{\overleftrightarrow{\mathcal{C}}}, \Sigma^+_{\#\$}, \overleftrightarrow{[\varepsilon]}, \delta_{\overleftarrow{\mathcal{C}}}, \delta_{\overrightarrow{\mathcal{C}}}, F_{\overleftrightarrow{\mathcal{C}}}).$$

Linear description of SCDAWGs. Our next goal is to show how to represent $\overleftrightarrow{\mathcal{C}}(\#\mathcal{D}\$)$ in space linear in the lexicon.

Proposition 5.11 *The following inequalities hold for the number of states $|Q_{\overleftrightarrow{\mathcal{C}}}|$, the number of the transitions in $\delta_{\overleftarrow{\mathcal{C}}}$, $|\delta_{\overleftarrow{\mathcal{C}}}|$, and the number of the transitions in $\delta_{\overrightarrow{\mathcal{C}}}$, $|\delta_{\overrightarrow{\mathcal{C}}}|$, in the SCDAWG $\overleftrightarrow{\mathcal{C}}(\#\mathcal{D}\$)$:*

$$|Q_{\overleftrightarrow{\mathcal{C}}}| \leq 2||\#\mathcal{D}||,$$

$$max(|\delta_{\overleftarrow{\mathcal{C}}}|, |\delta_{\overrightarrow{\mathcal{C}}}|) \leq 2||\#\mathcal{D}|| - 1.$$

Proof. First we shall give upper bounds for size of the *suffix tree* for $\#\mathcal{D}\$$, defined as the generalized deterministic finite-state automaton $STree(\#\mathcal{D}\$) := (Q, \Sigma_{\#\$}, \overrightarrow{\varepsilon}, \delta, F)$ where $Q := \{\overrightarrow{X} | X \in Subs(\#\mathcal{D}\$)\}$, $F := Suf(\#\mathcal{D}\$)$ and for $\sigma \in \Sigma_{\#\$}$ and $U \in \Sigma^+_{\#\$}$, $\delta(\overrightarrow{X}, \sigma U) = \overrightarrow{\overrightarrow{X} \circ \sigma}$ iff $\overrightarrow{X} \circ \sigma U \in Q$ and $\overrightarrow{\overrightarrow{X} \circ \sigma} = \overrightarrow{X} \circ \sigma U$, [IHS+01]. The suffix tree represents a tree with root $\overrightarrow{\varepsilon}$ and leaves F. For the number of the leaves we have $|F| \leq 1 + ||\#\mathcal{D}||$. Each internal node $\overrightarrow{X}$ of the suffix tree has at least two successors. Hence $|Q| \leq 2||\#\mathcal{D}||$. For the number of transitions in the suffix tree we have $|\delta| \leq 2||\#\mathcal{D}|| - 1$. For every $X \in Subs(\#\mathcal{D}\$)$ it can be shown that $\overleftrightarrow{X} \in Q$ and for every transition $\overrightarrow{[X]} \stackrel{\sigma U}{\to} \overleftrightarrow{[Y]}$ in $\delta_{\overleftarrow{\mathcal{C}}}$ it can be shown that there is a suffix tree transition $\overleftrightarrow{X} \stackrel{\sigma U}{\to} \overleftrightarrow{X\sigma U}$ in δ. Consequently $|Q_{\overleftrightarrow{\mathcal{C}}}| \leq |Q|$ and $|\delta_{\overleftarrow{\mathcal{C}}}| \leq |\delta|$. Analogously we obtain that $|\delta_{\overrightarrow{\mathcal{C}}}|$ is bounded by the number of transitions in $STree(\$\mathcal{D}^{rev}\#)$. □

[2] A state q is called *implicit* iff q is not the start state, q is not final and q has exactly one outgoing transition. A state is called *explicit* iff it is not implicit.

Remark 5.12 Note that Proposition 5.11 is not sufficient to prove that the size of SCDAWG is $O(||\#\mathcal{D}\$||)$, since the labels of the transitions are strings in $\Sigma^+_{\#\$}$. To achieve a linear description of SCDAWG the transitions are represented as follows. Let D be a concatenation of all strings in $\#\mathcal{D}\$$. For every state $q = \overleftrightarrow{[Y]}$ we store a position $end(q) = i$ in D where $\overleftrightarrow{Y}$ terminates, $\overleftrightarrow{Y} = D_{i-|\overleftrightarrow{Y}|+1} D_{i-|\overleftrightarrow{Y}|+2} \dots D_i$. For every transition $t = \overleftrightarrow{[X]} \xrightarrow{\alpha} \overleftrightarrow{[Y]}$ in $\delta_{\overleftarrow{\mathcal{C}}}$ let $start(t) := end(\overleftrightarrow{[Y]}) - |\alpha| + 1$. Since $\overleftrightarrow{X}\alpha$ is suffix of $\overleftrightarrow{Y}$, for every transition t in $\delta_{\overleftarrow{\mathcal{C}}}$ we store only $start(t)$, but not the whole label α. In analogous way we define $start(q)$ and $end(t)$ and for every transition t in $\delta_{\overrightarrow{\mathcal{C}}}$ we store only $end(t)$.

Online construction of SCDAWGs in linear time. In [IHS+01] Inenaga et al. present an online algorithm that builds letter by letter the SCDAWG $\overleftrightarrow{\mathcal{C}}(\{\#W\$\})$ for a single string $\#W\$$ in time $O(|\#W\$|)$. Here we present a new straightforward online algorithm that builds a representation of $\overleftrightarrow{\mathcal{C}}(\#\mathcal{D}\$)$ string by string in time $O(||\#\mathcal{D}\$||)$. Our result is essentially based on another algorithm by Inenaga et al. [IHS+05], that constructs $\overleftarrow{\mathcal{C}}(|\#\mathcal{D}\$|)$ letter by letter in an online manner. The idea is to synchronize $\overleftarrow{\mathcal{C}}(\#\mathcal{D}\$)$ and $\overleftarrow{\mathcal{C}}(\$\mathcal{D}^{rev}\#)$ while simultaneously building both of them word by word.

Proposition 5.13 *$\overrightarrow{\mathcal{C}}(\#\mathcal{D}\$)$ and $\overleftarrow{\mathcal{C}}(\$\mathcal{D}^{rev}\#)$ are isomorphic.*

Proof. Let $\overrightarrow{\mathcal{C}}(\#\mathcal{D}\$) = (Q, \Sigma_{\#\$}, \overleftrightarrow{[\varepsilon]}, \delta_{\overrightarrow{\mathcal{C}}}, F)$ and $\overleftarrow{\mathcal{C}}(\$\mathcal{D}^{rev}\#) = (Q', \Sigma_{\#\$}, \overleftrightarrow{[\varepsilon]}, \delta', F')$. The isomorphism is given by the bijection $b : Q \to Q'$ defined as follows.

$$b(\overleftrightarrow{[X]}) := \overleftrightarrow{[X^{rev}]}.$$

□

Since $\overleftarrow{\mathcal{C}}(\#\mathcal{D}\$)$ and $\overrightarrow{\mathcal{C}}(\#\mathcal{D}\$)$ have one and the same set of states, the bijection b, defined in the above proof, provides the way to express $\delta_{\overrightarrow{\mathcal{C}}}$ as follows:

$$\delta_{\overrightarrow{\mathcal{C}}}(q, \alpha) = b^{-1}(\delta'(b(q), \alpha)).$$

In our online construction we compute $\delta_{\overleftarrow{\mathcal{C}}}$, δ' and all values of b and b^{-1} for every state q. Let us note that if we directly compute $b(\overleftrightarrow{[X]})$ for a given state $\overleftrightarrow{[X]} \in Q$ by reversing some $Y \in \overleftrightarrow{[X]}$ and traversing $\overleftarrow{\mathcal{C}}(\$\mathcal{D}^{rev}\#)$ with Y^{rev} from the initial state, the total time for the whole construction would be in the worst case quadratic w.r.t. $||\#\mathcal{D}\$||$. To achieve linear time we need to compute $b(\overleftrightarrow{[X]})$, given a state $\overleftrightarrow{[X]}$, in amortized time $O(1)$. We show that such an efficient online computation of b can be based on the *suffix links* provided for every state $\overleftrightarrow{[X]}$ during the construction of $\overleftarrow{\mathcal{C}}(\#\mathcal{D}\$)$.

Definition 5.14 The *suffix link* $sl(\overleftrightarrow{[X]})$ of a state $\overleftrightarrow{[X]}$ in $\overleftarrow{\mathcal{C}}(\#\mathcal{D}\$)$ is $\overleftrightarrow{[Y]}$ where $\overleftrightarrow{Y}$ is the longest suffix of $\overleftrightarrow{X}$ such that $\overleftrightarrow{Y} \notin \overleftrightarrow{[X]}$. $sl(\overleftrightarrow{[\varepsilon]})$ is not defined.

Proposition 5.15 *Let*

$$\begin{aligned} \overleftarrow{\mathcal{C}}(\#\mathcal{D}\$) &= (Q, \Sigma_{\#\$}, \overleftrightarrow{[\varepsilon]}, \delta_{\overleftarrow{\mathcal{C}}}, F), \\ \overleftarrow{\mathcal{C}}(\$\mathcal{D}^{rev}\#) &= (Q', \Sigma_{\#\$}, \overleftrightarrow{[\varepsilon]}, \delta', F'), \end{aligned}$$

$\overleftrightarrow{[X]} \in Q$, $sl(\overleftrightarrow{[X]}) = \overleftrightarrow{[Y]}$ *and* $b(\overleftrightarrow{[Y]}) = q'$. *Let* $j = |\overleftrightarrow{X}| - |\overleftrightarrow{Y}|$, $\sigma = \overleftrightarrow{X}_j$. *Then there exists a* σ*-transition from* q'. *Let* $\delta'(q', \sigma U) = p'$ *be the* σ*-transition from* q'. *Then* $b(\overleftrightarrow{[X]}) = p'$.

Proof. The σ-transition from $q' = \overleftrightarrow{[Y^{rev}]}$ is defined, since $\overleftrightarrow{Y^{rev}}$ is prefix of $\overleftrightarrow{X^{rev}}$, which implies that there is a path with label $\overleftrightarrow{X}_j \overleftrightarrow{X}_{j-1} \ldots \overleftrightarrow{X}_1$ in $\overleftarrow{\mathcal{C}}(\$\mathcal{D}^{rev}\#)$ from q' to $\overleftrightarrow{[X^{rev}]}$. If we assume that this path is not composed of one single transition, then for the last intermediate state $\overleftrightarrow{[Z^{rev}]}$ of this path we have that $\overleftrightarrow{Z}$ is suffix of $\overleftrightarrow{X}$, $\overleftrightarrow{Z}$ is longer than $\overleftrightarrow{Y}$ and $\overleftrightarrow{Z} \notin \overleftrightarrow{[X]}$, which contradicts with $sl(\overleftrightarrow{[X]}) = \overleftrightarrow{[Y]}$. □

The algorithm of Figure 3 calculates $\overleftarrow{\mathcal{C}}(\#\mathcal{D}\$)$, $\overleftarrow{\mathcal{C}}(\$\mathcal{D}^{rev}\#)$, b and b^{-1}, given the lexicon $\#\mathcal{D}\$$. The states of these two CDAWGs are consecutive integers starting from 0, which is the initial state. The function $AddStringInCDAWG(\overleftarrow{\mathcal{C}}, \#W\$)$ represents the online construction of CDAWG invented by Inenaga et al., [IHS+05]. $AddStringInCDAWG$ adds the string $\#W\$$ to $\overleftarrow{\mathcal{C}}$. $AddStringInCDAWG$ changes its first argument $\overleftarrow{\mathcal{C}}$ by adding new consecutive states in Q and F and by setting the transition function $\delta_{\overleftarrow{\mathcal{C}}}$ and the suffix links sl for every new state. $AddStringInCDAWG$ never changes $\overleftrightarrow{X}$ for every state $\overleftrightarrow{[X]}$ that is already in Q. Hence the computation of b is stable in the sense that once $b(q)$ is precomputed for a given state q, further changes of $b(q)$ are impossible. Based on Proposition 5.15 the function $FindState$ recursively calculates the values of b. In line 5 of $FindState$ we use $end(i)$, the concatenation D of the strings accumulated so far described in Remark 5.12 and the *length* of state s defined as the length of the longest member of the equivalence class represented by s. The lengths of the states are computed by $AddStringInCDAWG$. The bottom of the recursion is guaranteed by $b(0) = 0$ and the decreasing lengths of the input states provided in recursive calls. The number of times $FindState$ is invoked is $O(|Q|)$. Since the time for the online construction of CDAWG is $O(||\#\mathcal{D}\$||)$ we obtain the following.

Proposition 5.16 *The online algorithm on Figure 3 runs in time* $O(||\#\mathcal{D}\$||)$.

The SCDAWG for $\#\mathcal{D}\$$ can be considered as a compact version of the bidirectional suffix trie $BiSTrie(\#\mathcal{D}\$)$, Definition 2.5.

Example 5.17 The SCDAWG for the example lexicon

$$\#\mathcal{D}\$ = \{\#\texttt{ear}\$, \#\texttt{lead}\$, \#\texttt{real}\$\}$$

```
BuildSDAWG( #D$ ){
1     Q = {0}; δ = ∅; F = ∅; Q' = {0}; δ' = ∅; F' = ∅;
2     ←C = (Q, Σ_#$, 0, δ, F); ←C' = (Q', Σ_#$, 0, δ', F'); b = ∅; b^-1 = ∅;
3     b(0) = 0; b^-1(0) = 0;
4     for( #W$ ∈ #D$ ){
5         n = |Q|;
6         AddStringInCDAWG(←C, #W$);
7         AddStringInCDAWG(←C', $W^rev#);
8         i = n;
9         while( i < |Q| ){
10            b(i) = nil; i = i + 1;
11        }
12        i = n;
13        while( i < |Q| ){
14            if( b(i) == nil ){
15                b(i) = FindState(←C, ←C', b, b^-1, i); b^-1(b(i)) = i;
16            }
17            i = i + 1;
18        }
19    }
20    return (←C, ←C', b, b^-1);
}
FindState( ←C, ←C', b, b^-1, i ){
1     s = sl(i);
2     if( b(s) == nil ){
3         b(s) = FindState(←C, ←C', b, b^-1, s); b^-1(b(s)) = s;
4     }
5     q' = b(s); j = end(i) − length(s); σ = D_j;
6     let δ'(q', σU) = p' be the σ-transition from q' in ←C'
7     return p';
}
```

Figure 3: Online construction of a representation of SCDAWG for $\#\mathcal{D}\$$ as $(\overleftarrow{C}\ (\#\mathcal{D}\$), \overleftarrow{C}\ (\$\mathcal{D}^{rev}\#), b, b^{-1})$

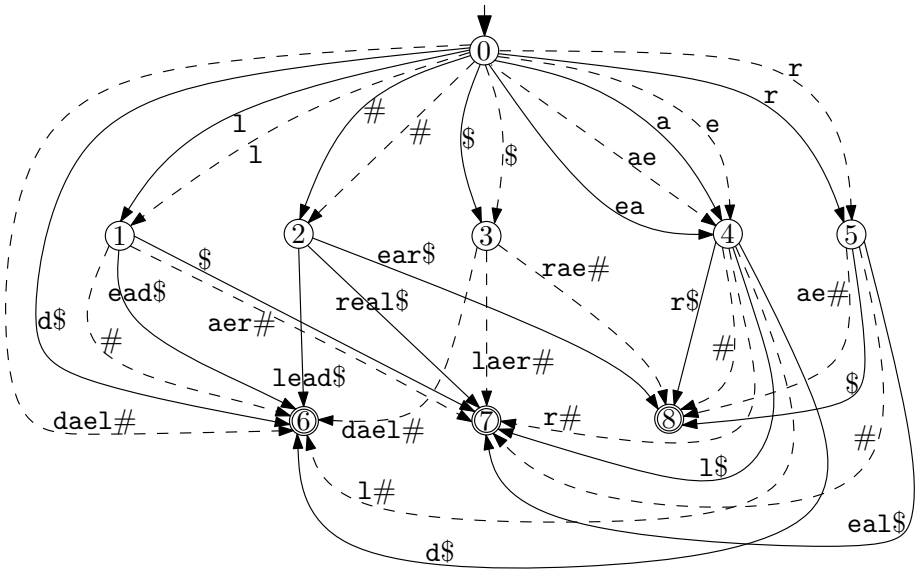


Figure 4: SCDAWG for $\{\#\texttt{ear\$}, \#\texttt{lead\$}, \#\texttt{real\$}\}$.

is shown in Figure 4. The dashed transitions represent $\delta_{\overrightarrow{\mathcal{C}}}$, while the solid transitions represent $\delta_{\overleftarrow{\mathcal{C}}}$. The equivalence classes are 0 - $\{\varepsilon\}$, 1 - $\{\texttt{l}\}$, 2 - $\{\#\}$, 3 - $\{\$\}$, 4 - $\{\texttt{a}, \texttt{ea}\}$, 5 - $\{\texttt{r}\}$, 6 - $\{\texttt{d\$}, \texttt{ad\$}, \texttt{lead\$}, \#\texttt{lead\$}\}$, 7 - $\{\texttt{l\$}, \texttt{al\$}, \texttt{eal\$}, \texttt{real\$}, \#\texttt{real\$}\}$ and 8 - $\{\texttt{r\$}, \texttt{ar\$}, \texttt{ear\$}, \#\texttt{ear\$}\}$.

5.2 Bidirectional online search using SCDWAGs

We now describe how the above index structure is used for computation of solution sets defined in Section 4.2. We assume that the following **offline resources** are available:

1. the SCDAWG $\overleftrightarrow{\mathcal{C}}(\#\mathcal{D}\$)$, in particular the two transition funtions $\delta_{\overleftarrow{\mathcal{C}}}$ and $\delta_{\overrightarrow{\mathcal{C}}}$;
2. $start(\overleftrightarrow{[V]})$ and $end(\overleftrightarrow{[V]})$ for every state $\overleftrightarrow{[V]}$, Remark 5.12;
3. $start(t)$ for every transition in $\delta_{\overleftarrow{\mathcal{C}}}$, Remark 5.12;
4. $end(t)$ for every transition in $\delta_{\overrightarrow{\mathcal{C}}}$, Remark 5.12.

We keep track of the following **online information** - here W denotes the substring of a lexicon word faced at a certain point of the computation of solution sets and D denotes the concatenation used in the linear representation of the SCDAWG $\overleftrightarrow{\mathcal{C}}(\#\mathcal{D}\$)$, 5.12.

1. the length of W;

2. (the number of) the state $\overleftrightarrow{[W]}$;

3. the unique (Proposition 5.4) position j_W of W in D such that $start(\overleftrightarrow{[W]}) \leq j_W \leq end(\overleftrightarrow{[W]})$.

Let $\sigma \in \Sigma_{\#\$}$. We consider possible extensions of the current substring W to the right of the form $W\sigma$ as follows. If $j_W+|W| \leq end(\overleftrightarrow{[W]})$, then $W\sigma \in Subs(\#\mathcal{D}\$)$ iff $\sigma = D_{j_W+|W|}$ and if $\sigma = D_{j_W+|W|}$, then $\overleftrightarrow{[W\sigma]} = \overleftrightarrow{[W]}$ and $j_{W\sigma} = j_W$. If $j_W + |W| > end(\overleftrightarrow{[W]})$, then $W\sigma \in Subs(\#\mathcal{D}\$)$ iff there exists a σ-transition from $\overleftrightarrow{[W]}$ in $\delta_{\overleftarrow{\mathcal{C}}}$. Let $t = \overleftrightarrow{[W]} \xrightarrow{\sigma U} \overleftrightarrow{[V]}$ be the σ-transition from $\overleftrightarrow{[W]}$ in $\delta_{\overleftarrow{\mathcal{C}}}$. Then $\overleftrightarrow{[W\sigma]} = \overleftrightarrow{[V]}$ and $j_{W\sigma} = start(t) - |W|$. Possible extensions to the left of the form σW are handled similarly by using $start(\overleftrightarrow{[W]})$ and $end(t)$.

Example 5.18 One example for the use of $\overleftrightarrow{\mathcal{C}}(\#\mathcal{D}\$)$ in Figure 4 is the following. We first want to check if `e` is a substring in $Subs(\#\mathcal{D}\$)$. For this aim we start from state 0 and follow the `e`-transition $t = 0 \xrightarrow{\texttt{ea}} 4$ in $\delta_{\overleftarrow{\mathcal{C}}}$, $j_{\texttt{e}} = start(t) = start(4) = end(4) - 1$, the number of the state $\overleftrightarrow{[\texttt{e}]}$ is 4.

- We now want to find all left extensions with a single letter. Since $j_{\texttt{e}} = start(4)$, we have to use the three possible dashed transitions from state 4. With # we reach 8, the number of the state $\overleftrightarrow{[\#\texttt{e}]}$ is 8, $j_{\#\texttt{e}} = start(8) = end(8) - 4$. With `r` we reach 7, the number of the state $\overleftrightarrow{[\texttt{re}]}$ is 7, $j_{\texttt{re}} = start(7) + 1 = end(7) - 4$. With `l` we reach 6, the number of the state $\overleftrightarrow{[\texttt{le}]}$ is 6, $j_{\texttt{le}} = start(6) + 1 = end(6) - 4$.

- We now want to find all right extensions of `e` with a single letter. Since $j_{\texttt{e}} + 1 = end(4)$, the only one possible extension to the right is with letter `a`, the number of the state $\overleftrightarrow{[\texttt{ea}]}$ is 4, $j_{\texttt{ea}} = j_{\texttt{e}} = start(4) = end(4) - 1$. If we want to further extend `ea` to the right, we have to use the solid transition, because $j_{\texttt{ea}} + 2 > end(4)$.

Remark 5.19 In our actual implementation we use a simple optimization of the approximate search based on additional information stored in the SCDAWG. The idea is to use "positional" information to recognize blind paths of the search. Consider a substring V of a lexicon word. If some lexicon word W has the form $W = RVS$ we say that $|R|$ (resp. $|S|$) is the *length of a possible prefix (suffix) for* V. In the SCDAWG we store for each substring $\overleftrightarrow{V}$ the maximal and minimal length of a possible prefix (suffix) for $\overleftrightarrow{V}$. When computing solution sets $Sol_{\mathcal{D}}(P', b')$, substrings P_0 of the pattern P are aligned with substrings V found in the SCDAWG. Each substring P_0 defines a unique prefix R_P and a unique suffix S_P of the pattern $P = R_P P_0 S_P$. We check if the length of R_P and S_P is "compatible" with the information stored in the SCDAWG for the length of possible prefixes and suffixes for V. To test "compatibility", the error bound

and the distance between P_0 and V has to be taken into account. Compatibility is checked each time we reach new state of the SCDAWG. We omit the technical details.

6 Evaluation

In this section we compare our new method to two other methods for efficient approximate search in lexica, Oflazer's approach [Ofl96] and the forward-backward method introduced in [MS04]. To have a common basis for the experiments we always use as a filter mechanism Ukonnen's optimized matrix method [Ukk85].[3] For the three methods we present experimental results for approximate search in lexica of different sizes and types. We also look at the dependency of search times on the notion of similarity used. In order to get a picture of principle limitations for approximate search we also present evaluation results where we simulate the "ideal method".

The "ideal method" for bound $b \in \mathbb{N}$ and dictionary $\mathcal{D}$ is based on a *perfect index* $I(\mathcal{D}, b)$ that directly maps every query (P, b) to the solution $Sol_{\mathcal{D}}(P, b) \cap \mathcal{D}$. Since the size of the perfect index $I(\mathcal{D}, b)$ would be too large, for every experiment we build a *restricted perfect index* $I(T, \mathcal{D}, b)$ that works only for a small finite test set $\mathcal{T} \subset \Sigma^*$ of query strings. For every query string $P \in \mathcal{T}$ the restricted perfect index $I(\mathcal{T}, \mathcal{D}, b)$ maps (P, b) to the solution $Sol_{\mathcal{D}}(P, b) \cap \mathcal{D}$. We represent the restricted perfect index $I(\mathcal{T}, \mathcal{D}, b)$ as an acyclic k-subsequential transducer for $k = max_{P \in \mathcal{T}} |Sol_{\mathcal{D}}(P, b) \cap \mathcal{D}|$. An online algorithm for building minimal acyclic k-subsequential transducers is introduced in [MM01]. This form of representation is optimal since the only time used is the time for reading the input and directly producing the desired output.

6.1 Comparison of search times for different methods

For our first series of experiments we chose a lexicon $\mathcal{D}$ of $1,200,070$ book titles. The average length of titles is 47.64. The number of different symbols in the alphabet of the lexicon is 99. We compare search times obtained for Oflazer's method [Ofl96], the forward-backward method [MS04], the new method and the "ideal method". In all experiments we set the weight of each nonidentity edit operation op to $w(op) = 1$. We then vary the distance bound b from 2 to 15. For each bound b we generated a test set $\mathcal{T}$ of $10,000$ query strings. Each query string P was received from a randomly chosen string $W \in \mathcal{D}$ by applying randomly b operations from the set of edit operations Op to W such that $|P| \geq 3b$.

All experiments were run on a machine with 64 gygabytes of RAM, two 2.4 GHz Quad-Core Intel Xeon 8-core processors, 256 KB L2 cache memory per core and 12 MB L3 cache memory per processor. Our implementation uses only

[3] Universal Levenshtein automata [SM02, MS04] are more efficient but can only be built for small distance bounds because of huge memory requirements.

Levenshtein distance, lexicon of book titles				
b	new / ideal	fb / ideal	f / ideal	ideal (ms)
2	13.87	325.083	3817.47	0.004
3	22.80	686.037	20063.82	0.004
4	50.72	3904.30		0.004
5	54.24	7741.55		0.005
6	76.36			0.005
7	86.55			0.005
8	173.98			0.006
9	154.32			0.006
10	172.49			0.006
11	163.06			0.007
12	287.94			0.007
13	261.13			0.007
14	301.03			0.008
15	300.05			0.008

Table 1: Comparison of search times for four different methods, standard Levenshtein distance, dictionary of titles. f - Oflazer's method, fb - the forward-backward method, new - the new method, ideal - the "ideal method". Empty cells mean that we did not wait for termination. Explicit search times (in ms) are only given for the ideal method (last column). All other entries represent factors, comparing the given method with the ideal method.

one thread. The amount of memory needed for our experiments is determined by the size of the precomputed index[4].

Table 1 presents results obtained for the standard Levenshtein distance. Column 1 specifies the value of the distance bound b used in the experiments. Explicit search times are only presented for the ideal method (column 5, times in milliseconds). Numbers x in Table 1 for some method M mean that the ideal method was x times faster than method M for the problem class. For example, the entry 13.87 found in row/column 2 indicates that approximate search using the new method presented above with distance bound 2 and standard Levenshtein distance on average took 13.87 times the time needed by the ideal method. Here, as in all experiments, the time needed to write the output words is always included. Empty cells found in the table mean that we did not wait for the respective method to finish.

The results in Table 1 show that the method presented in this paper comes "close" to the ideal method for small distance bounds when using the standard Levenshtein distance. For the given lexicon of titles, which contains long strings, the new method is dramatically faster than the forward-backward method, which in turn is much faster than Oflazer's method. It is worth to note that the differences become more and more drastic when using larger distance bounds. For these bounds only the new method leads to acceptable search times.

[4]We use depth first implementations of the evaluated methods, see Remark 4.11.

b	Levenshtein distance, lexicon of MEDLINE sentences		
	new / ideal	fb / ideal	f / ideal
2	8.64	71.62	664.91
3	10.72	145.60	2670.98
4	13.55	727.22	
5	16.57	1357.87	
6	24.54		
7	27.23		
8	36.47		
9	40.68		
10	59.48		
11	62.36		
12	123.31		
13	123.84		
14	145.98		
15	146.11		
20	411.71		
30	1552.28		
40	4872.49		
50	23869.91		

Table 2: Comparison of search times for different methods, standard Levenshtein distance, dictionary of MEDLINE sentences. f - Oflazer's method, fb - the forward-backward method, new - the new method, ideal - the "ideal method". All entries represent factors, comparing the given method with the ideal method. Empty cells mean that we did not wait for termination.

6.2 Comparison of search times for language databases with sentences

For our second series of experiments we use a collection of sentences from the life sciences and biomedical domain. The lexicon consists of all sentences from $43,000$ paragraphs which were randomly chosen from MEDLINE abstracts[5]. The number of sentences in our list is $351,008$. The average number of symbols per sentence is 149.26. The size of the lexicon is approximately the same as the size of the lexicon of titles, but the strings are longer. Table 2 presents the comparison of the different methods for the standard Levenshtein distance. As a new challenge, the distance bound b used for approximate search varies from 2 to 50. Note that for previous methods the use of larger distance bounds leads to unacceptable search times. Speed-up factors are similar to those observed in Table 1.

6.3 Comparison of search times for different variants of Levenshtein distance

For our third series of experiments we compare search times obtained for three notions of similarity, (i) the standard Levenshtein distance, (ii) the variant where transpositions of neighbored symbols are treated as additional edit operations,

[5]MEDLINE is a bibliographic database of U.S. National Library of Medicine. MEDLINE contains over 19 million references to journal articles in life sciences, www.nlm.nih.gov/pubs/factsheets/medline.html.

Levenshtein distance with transpositions, lexicon of book titles			
b	new / ideal	fb / ideal	f / ideal
2	19.69	346.28	4081.83
3	36.00	1022.59	20536.91
4	93.36	4767.08	
5	108.54	11861.07	
6	149.84		
7	167.16		
8	315.19		
9	271.12		
10	324.32		
11	310.59		
12	467.27		
13	459.23		
14	500.79		
15	496.84		

Table 3: Comparison of search times for the variant of Levenshtein distance transpositions of neighbored symbols are treated as edit operations. f - Oflazer's method, fb - the forward-backward method, new - the new method, ideal - the "ideal method". All entries represent factors, comparing the given method with the ideal method. Empty cells mean that we did not wait for termination.

and (iii) the variant where also merges and splits are used as additional edit operations. Tables 3 (resp. Table 4) presents results obtained for the variant of Levenshtein distance where we also use transpositions (merges and splits) as operations.

Basically, the results in Tables 3 and 4 are similar to the previous results: also for the modified distances, the new method is much faster than previous methods, with a speed-up factor of at least 15. For large distance bounds the speed-up factor is larger.

6.4 Comparison of search times for different symbol distributions

In our fourth series of experiments we ask how the statistical properties of the distribution of letters in the lexicon words influence search times. We generated two random dictionaries of $1,200,070$ strings - one with uniform distribution of 99 symbols and average string length 54.35 and another one with binomial distribution of 99 symbols and average string length 54.63. In Table 5 the columns "Lev+Binomial" and "Lev+Uniform" present the behavior of the algorithms for the two random dictionaries with binomial and uniform distributions of the symbols.

The differences between the search times for the three methods for approximate search observed in Table 1 basically remain unchanged. For bound $b = 2$, search in the natural language lexicon of titles (Table 1) is faster than search in the lexica with binomial distribution and is slower than search in the lexica with uniform distribution. For larger distance bounds, the differences between the search times for the three types of lexica are more difficult to interpret.

Levenshtein distance with merges and splits, lexicon of book titles			
b	new / ideal	fb / ideal	f / ideal
2	51.69	2297.07	37428.91
3	155.55	7199.24	
4	693.20	42141.64	
5	788.27	112676.78	
6	1091.35		
7	1229.88		
8	2267.31		
9	1868.38		
10	2462.66		
11	2151.49		
12	3352.05		
13	3077.23		
14	3350.26		
15	3210.70		

Table 4: Comparison of search times for the variant of Levenshtein distance where merges of two symbols into one and splits of a symbol into two symbols are treated as edit operations. f - Oflazer's method, fb - the forward-backward method, new - the new method, ideal - the "ideal method". All entries represent factors, comparing the given method with the ideal method. Empty cells mean that we did not wait for termination.

b	Lev+Binomial			Lev+Uniform		
	new / ideal	fb / ideal	f / ideal	new / ideal	fb / ideal	f / ideal
2	18.12	435.29	12513.00	10.68	1686.54	74164.46
3	20.12	861.60	103772.38	12.14	3392.04	262620.28
4	22.04	11936.27		14.59	73079.20	
5	25.30	23053.83		17.33	136324.69	
6	38.04			25.22		
7	41.01			27.06		
8	53.56			33.69		
9	60.33			39.79		
10	81.73			56.57		
11	89.99			61.09		
12	156.64			114.70		
13	164.48			124.96		
14	188.93			137.76		
15	188.52			136.48		

Table 5: Comparison of search times for different symbol distributions. Randomly generated lexica. f - Oflazer's method, fb - the forward-backward method, new - the new method, ideal - the "ideal method". All entries represent factors, comparing the given method with the ideal method. Empty cells mean that we did not wait for termination.

b	Levenshtein distance, lexicon of Bulgarian word forms		
	new / ideal	fb / ideal	f / ideal
2	85.86	207.10	1980.55
3	105.28	510.04	9236.49
4	420.32	1905.09	

Table 6: Search times for Bulgarian lexicon with short strings showing the influence of the length of lexicon entries. Standard Levenshtein distance, search times for three methods. f - Oflazer's method, fb - the forward-backward method, new - the new method, ideal - the "ideal method".

method	new	fb	f
Titles	1104.63	110.32	54.18
MEDLINE	921.41	100.02	49.87
Bg word forms	61.02	9.22	4.17

Table 7: Dictionary of titles and dictionary of MEDLINE sentences vs. smaller dictionary, sizes of indexes in megabytes.

6.5 Influence of the length of the strings in the lexicon

In our last experiment we look at the influence of the length of the strings in the lexicon. We selected a smaller dictionary of natural language expressions consisting of approximately $450,000$ Bulgarian word forms with average word length 10.01. In Table 6 the corresponding search times for the smaller dictionary are found in columns $2-4$. Since for every query string P we require $|P| \geq 3b$, for bounds $b > 4$ there are less than $10,000$ entries in the smaller dictionary from which we could generate queries. For this reason in the case of the smaller dictionary we do not present results for $b > 4$. Even for the short strings of the Bulgarian lexicon, the new method is much faster than the forward-backward method and the third method. The speed-up gained is less drastic than for the lexicons of titles and MEDLINE sentences, and here the "ideal method" remains more than 85 times faster than the new method.

Size of index structures. Table 7 represents for every method, except the ideal one, the sizes in megabytes of the indexes compiled from the dictionary of titles, the dictionary of MEDLINE sentences and the dictionary of Bulgarian word forms.

7 Historical remarks, possible applications and conclusion

We introduced a new method for fast approximate search in lexica that can be used for a large family of string distances. The method uses a bidirectional index structure for the lexicon. This index structure can be seen as a part of a longer development of related index structures starting with work on suffix tries, suffix trees, and directed acyclic word graphs (DAWGs) [Wei73, McC76,

Ukk95, CS84, BBH+85]. These index structures address single texts and are one-directional in the sense that search for substrings of the given string/text follows the left-to-right reading order. In [Sto95, Sto00, Maa00, IHS+01] it has been shown how to obtain *bidirectional* index structures for strings/texts, supporting search for substrings using both left-to-right and right-to-left reading order. One-directional index structures for *sets of strings* (as opposed to single strings) have been described in [BBH+87, Gus97, Bre98, MMW09, IHS+05]. In each case the challenge is to find an index structure with size linear in the size of the input text or lexicon, with a linear-time construction algorithm. In [BBH+87] a *bidirectional index structure for sets of texts* is briefly sketched, asking for natural applications. In this paper we have seen that such an index applied to lexica can be used to realize a very fast method for approximate search.

With the new index, the "wall effect" mentioned in the Introduction can be avoided. Among related techniques, the BLASTA method [AGM+90] is worth mentioning. In this approach, the occurrences of specific substrings in the lexicon are indexed in order to reduce the lexicon words to be considered. It assumes that each answer of the query has to contain at least one of the keyed substrings which allows it to start with an exact match of such a promising substring. In such a way BLASTA prunes the initial exhaustive search and proves to be efficient. However since there is no guarantee that all answers of the query meet this condition, it may fail to retrieve the complete list of words satisfying the query.

Our evaluation results show that the new method is much faster than previous methods, and for lexica with long strings the speed-up is drastic. Here the new method for distance bound $b = 2$ comes close to the theoretical limit when using the standard Levenshtein distance.

We add a brief comment on possible applications. As a matter of fact, the method may be used to speed-up traditional spelling correction techniques. For high quality spelling correction, speed is not the only issue. Current approaches typically use probabilistic techniques at two places. First, good similarity measures for selecting candidates are based on special edit operations with weights depending on the particular symbols/strings used. How to find appropriate edit operations and weights is a question beyond the scope of the current paper. However, the framework of a generalized distance we use to model similarities should be general enough to cover most interesting cases. Second, when looking for an optimal correction suggestion for a misspelled token, language models (e.g., weighted word trigrams) help to find a correction suggestion that fits the local context. Still, *similarity search* in the background lexicon only looks at *single* tokens and for efficiency reasons, "context sensitive" correction suggestions for distinct tokens are often computed in isolation. An interesting question is if better results are obtained when using larger contexts already for the background lexica and similarity search. This strategy would guarantee that the correction suggestions obtained for a sequence of tokens always fit together. The method introduced above offers new possibilities for testing such a strategy since we can use large strings and distance bounds. As a matter of fact, issues of

smoothing have to be taken into account when trying to synchronize contextual similarity search and language models.

Possible application areas of the new method are not restricted to traditional fields of approximate search such as spell checking, text and OCR correction. Since the method is fast enough to deal with collections of long strings and large distance bounds, it seems promising to test its use, e.g., for detecting plagiarism, for finding similar sentences in translation memories and related language databases, and for approximate search in collections of address or bibliographic data. We currently also look at a variant of the method for fast approximate search of patterns in an indexed collection of texts. In order to find all approximate matches for a string in a (collection of) texts, the index has to be enriched by adding information on the positions of all occurrences of each infix. The challenge is to keep the size of the index linear in the size of the text(s).

A remaining open question is the time complexity of the presented algorithm. A desirable approach would be to estimate the average complexity in a way similar to Myers' [Mye94].

There are two obvious ways how search times presented above could be immediately improved. First, for small distance bound we could use universal Levenshtein automata [MS04] as filters. This leads to a performance gain as compared to matrix based filters [MMS11]. Second, an additional speed-up could be obtained by running subsearches of distinct branches of the search tree used in parallel. The optimal selection of search trees is an interesting point for further investigations.

References

[AGM+90] Stephen F. Altschul, Warren Gish, Webb Miller, Eugene W. Myers, and David J. Lipman. A basic local alignment search tool. *Journal of Molecular Biology*, 215:403–410, 1990.

[AK05] Farooq Ahmad and Grzegorz Kondrak. Learning a spelling error model from search query logs. In *HLT '05: Proceedings of the Conference on Human Language Technology and Empirical Methods in Natural Language Processing*, pages 955–962, Vancouver, British Columbia, Canada, 2005.

[BBH+85] A. Blumer, J. Blumer, D. Haussler, A. Ehrenfeucht, M. T. Chen, and J. Seiferas. The smallest automation recognizing the subwords of a text. *Theoretical Computer Science*, 40:31 – 55, 1985. Eleventh International Colloquium on Automata, Languages and Programming.

[BBH+87] A. Blumer, J. Blumer, D. Haussler, R. McConnell, and A. Ehrenfeucht. Complete inverted files for efficient text retrieval and analy-

sis. *Journal of the Association for Computing Machinery*, 34(3):578–595, 1987.

[BCP02] Ilaria Bartolini, Paolo Ciaccia, and Marco Patella. String matching with metric trees using an approximate distance. In *String Processing and Information Retrieval*, pages 271–283, 2002.

[Bre98] Dany Breslauer. The suffix tree of a tree and minimizing sequential transducers. *Theoretical Computer Science*, 191(1-2):131 – 144, 1998.

[BYN98] Ricardo A. Baeza-Yates and Gonzalo Navarro. Fast approximate string matching in a dictionary. In *String Processing and Information Retrieval*, pages 14–22, 1998.

[BYN99] Ricardo A. Baeza-Yates and Gonzalo Navarro. Faster approximate string matching. *Algorithmica*, 23(2):127–158, 1999.

[CB04] Silviu Cucerzan and Eric Brill. Spelling correction as an iterative process that exploits the collective knowledge of web users. In *EMNLP '04: Proceedings of 2004 Conference on Empirical Methods in Natural Language Processing*, pages 293–300, 2004.

[CS84] M. T. Chen and J. Seiferas. Efficient and elegant subword tree construction. In A. Apostolico and Z. Galil, editors, *Combinatorial Algorithm on Words, volume 12 of NATO Advanced Science Institutes, Series F*, pages 97–107. Springer-Verlag, 1984.

[DHH+97] Andreas Dengel, Rainer Hoch, Frank Hönes, Thorsten Jäger, Michael Malburg, and Achim Weigel. Techniques for improving OCR results. In Horst Bunke and Patrick S.P. Wang, editors, *Handbook of Character Recognition and Document Image Analysis*, pages 227–258. World Scientific, 1997.

[GK97] Robert Giegerich and Stefan Kurtz. From Ukkonen to McCreight and Weiner: A Unifying View of Linear-Time Suffix Tree Construction. *Algorithmica*, 19:331–353, 1997.

[Gus97] Dan Gusfield. *Algorithms on Strings, Trees, and Sequences: Computer Science and Computational Biology.* Cambridge University Press, 1997.

[IHS+01] Shunsuke Inenaga, Hiromasa Hoshino, Ayumi Shinohara, Masayuki Takeda, and Setsuo Arikawa. On-line construction of symmetric compact directed acyclic word graphs. In *Proc. of 8th International Symposium on String Processing and Information Retrieval (SPIRE'01)*, pages 96–110. IEEE Computer Society, 2001.

[IHS+05] Shunsuke Inenaga, Hiromasa Hoshino, Ayumi Shinohara, Masayuki Takeda, Setsuo Arikawa, Giancarlo Mauri, and Giulio Pavesi. Online construction of compact directed acyclic word graphs. *Word Journal Of The International Linguistic Association*, 146(2):1–12, 2005.

[Kuk92] Karen Kukich. Techniques for automatically correcting words in texts. *ACM Computing Surveys*, pages 377–439, 1992.

[Lev66] Vladimir I. Levenshtein. Binary codes capable of correcting deletions, insertions, and reversals. *Soviet Physics Doklady*, 1966.

[LH99] Tessa Lau and Eric Horvitz. Patterns of search: analyzing and modeling web query refinement. In *UM '99: Proceedings of the Seventh International Conference on User Modeling*, pages 119–128, Secaucus, NJ, USA, 1999. Springer-Verlag New York, Inc.

[Maa00] Moritz G. Maass. Linear bidirectional on-line construction of affix trees. In *Proc. of 11th Ann. Symp. on Combinatorial Pattern Matching (LNCS1848)*, pages 320–334. Springer-Verlag, 2000.

[McC76] Edward M. McCreight. A space-economical suffix tree construction algorithm. *Journal of the Association for Computing Machinery*, 23(2):262–272, 1976.

[MM01] Stoyan Mihov and Denis Maurel. Direct construction of minimal acyclic subsequential transducers. In Sheng Yu and Andrei Paun, editors, *Implementation and Application of Automata, 5th International Conference, CIAA 2000, London, Ontario, Canada, July 24-25, 2000, Revised Papers*, volume 2088 of *Lecture Notes in Computer Science*, pages 217–229. Springer, 2001.

[MMS11] Petar Mitankin, Stoyan Mihov, and Klaus U. Schulz. Deciding word neighborhood with universal neighborhood automata. *Theoretical Computer Science*, 412(22):2340 – 2355, 2011.

[MMW09] Mehryar Mohri, Pedro Moreno, and Eugene Weinstein. General suffix automaton construction algorithm and space bounds. *Theoretical Computer Science*, 410(37):3553–3562, 2009.

[MS04] Stoyan Mihov and Klaus U. Schulz. Fast approximate search in large dictionaries. *Computational Linguistics*, 30(4):451–477, 2004.

[Mye94] Eugene W. Myers. A sublinear algorithm for approximate keyword searching. *Algorithmica*, 12:345–374, 1994.

[NBY99] Gonzalo Navarro and Ricardo A. Baeza-Yates. Very fast and simple approximate string matching. *Information Processing Letters*, 72:65–70, 1999.

[NBY00] Gonzalo Navarro and Ricardo Baeza-Yates. A hybrid indexing method for approximate string matching. *Journal of Discrete Algorithms*, 1(1):205–239, 2000.

[Ofl96] Kemal Oflazer. Error-tolerant finite-state recognition with applications to morphological analysis and spelling correction. *Computational Linguistics*, 22(1):73–89, 1996.

[SM02] Klaus U. Schulz and Stoyan Mihov. Fast string correction with Levenshtein automata. *IJDAR*, 5(1):67–85, 2002.

[Sto95] Jens Stoye. Affixbäume. Master's thesis, Universität Bielefeld, May 1995.

[Sto00] Jens Stoye. Affix trees. Technical Report 2000-04, Universität Bielefeld, Technische Fakultät, 2000.

[Ukk85] Esko Ukkonen. Algorithms for approximate string matching. *Information Control*, 64:100–18, 1985.

[Ukk95] Esko Ukkonen. On-line construction of suffix-trees. *Algorithmica*, 14(3):249–260, 1995.

[Ver88] Jean Veronis. Computerized correction of phonographic errors. *Computers and Humanities*, 22(1):43–56, 1988.

[Wei73] Peter Weiner. Linear pattern matching algorithms. In *Proceedings of 14th IEEE Annual Symposium on Switching and Automata Theory*, pages 1–11, 1973.

[WM92] Sun Wu and Udi Manber. Fast text searching: allowing errors. *Communications of the ACM*, 35(10):83–91, 1992.

Appendix

We show how the search strategy described in Section 4 can be adapted to the case of an arbitrary generalized distance $d = (Op, w)$. In what follows, $\omega_{\max}$ denotes the maximal width of an operation $op \in Op$. To simplify the following description, we introduce the notion of a (left, right) reduct of a word. Intuitively, reducts of a word U are obtained by deleting a "short" (possibly empty) prefix and/or suffix of length $< \omega_{max}$ from U.

Definition 7.1 Let $U \in \Sigma^*$ be represented in the form $U = U_1 \circ U_2$. If $|U_1| < \omega_{max}$, then U_2 is called a *left reduct* of U. If $|U_2| < \omega_{max}$, then U_1 is called a *right reduct* of U. If $U = U_1 \circ V \circ U_2$ and both $|U_1| < \omega_{max}$ and $|U_2| < \omega_{max}$, then V is called a *reduct* of U.

We denote that for $\omega_{max} = 1$ always U is the only reduct of U. The formal background for the adapted search procedure is provided by the following generalization of Proposition 4.7.

Proposition 7.2 *Let $P' = P'_1 \circ P'_2$ and α be an alignment with $l(\alpha) = P'$ and $w(\alpha) \leq b'$, then:*

1. *α can be represented in the form $\alpha = \alpha_1 \circ \beta \circ \alpha_2$ such that $\beta \in Op \cup \{\varepsilon\}$, $l(\alpha_1)$ is a right reduct of P'_1, and $l(\alpha_2)$ is a left reduct of P'_1.*

2. *for each such decomposition and integers b'_1 and b'_2 with $b'_1 + b'_2 = b' - 1$ it holds that $w(\alpha_1) \leq b'_1$ or $w(\alpha_2) \leq b'_2$.*

Proof. We first prove Part 1. Let α_1 denote the maximal prefix of α with the property that $l(\alpha_1)$ is a prefix of P'_1. If $l(\alpha_1) = P'_1$ we define $\beta := \varepsilon$. Otherwise there exists an operation $op \in Op$ such that $l(\alpha_1)$ is a proper prefix of P'_1, the latter being a proper prefix of $l(\alpha_1 \circ op)$. In this case we define $\beta := op$. In both cases α_2 is now determined by the equation $\alpha = \alpha_1 \circ \beta \circ \alpha_2$. It is trivial to check that this representation has the properties stated above. The second statement follows easily. □

Recall that in the special situation considered in Section 4 we decomposed the pattern P into subparts $P = P_1 \circ P_2 \cdots \circ P_{b+1}$, and for substrings of the form $P_k \circ \cdots \circ P_l$ ($k \leq l$, possible combinations of k, l determined by the structure of the search tree) we computed approximate matches with substrings of lexicon words using distinct bounds. In the general situation considered here we split P as above. We then try to find approximate matches between *reducts* of the substrings $P_k \circ \cdots \circ P_l$ with substrings of lexicon words. For a formal description, let us introduce another notational convention. By $r(i, U, j)$ we denote the reduct obtained from U by deleting the unique prefix and suffix of length i and j, respectively. Hence $r(0, U, 0) = U$.

Building the generalized search tree for a pattern. For a given input pattern P and a bound b, let $\mathcal{T}_P$ denote the search tree defined in Section 4. With each query (P', b') decorating a node η we associate as a subcase analysis the set of all *derived* queries of the form $(r(i, P', j), b')$ where $i, j < \omega_{max}$. The problem considered at node η is to solve all derived queries of the above form. Note that $(r(0, P, 0), b)$ is equivalent to (P, b).

Computation of solution sets for derived queries. For each derived query $(r(i, P', j), b')$ of the generalized tree we compute a set $S_D(r(i, P', j), b')$ in a bottom-up fashion. We shall prove below that $S_D(r(i, P', j), b')$ is the solution set $Sol_D(r(i, P', j), b')$ in each case.

Initialization steps. For a derived query $(r(i, P', j), 0)$ at a leaf we decide if $r(i, P', j)$ is a substring of a lexicon word. In the positive case we let $S_D(r(i, P', j), 0) := \{r(i, P', j)\}$, otherwise we define $S_D(r(i, P', j), 0) := \emptyset$.

Extension steps. Let $((r(i, P', j), b')$ denote a derived query at a non-leaf node η of $\mathcal{T}_P$, let let (P'_1, b'_1) and (P'_2, b'_2) denote the main queries of the two children η_1, η_2 of η, which are given in the natural left-to-right ordering. Given all sets $S_D((r(i, P'_1, j_1), b'_1)$ and $S_D((r(i_2, P'_2, j), b'_2)$ for the derived queries at η_1, η_2 we define $S_D(r(i, P', j), b')$ as the union of the two sets S_1 and S_2 defined

as

$$S_1 = \bigcup_{j_1=0}^{\omega_{max}-1} \{U \circ V \in Subs(\mathcal{D}) \mid U \in S_D(r(i, P'_1, j_1), b'_1), d^*_1 \leq b'\}$$

$$S_2 = \bigcup_{i_2=0}^{\omega_{max}-1} \{V \circ U \in Subs(\mathcal{D}) \mid U \in S_D(r(i_2, P'_2, j), b'_2), d^*_2 \leq b'\}\}.$$

Here $d^*_1 = d(r(i_1, P'_1, j_1), U) + d(Q_1, V)$ where Q_1 is obtained from $P' = P'_1P'_2$ by deleting the prefix of length $|P'_1| - j_1$ and the suffix of length j. Similarly $d^*_2 = d(r(i_2, P'_2, j), U) + d(Q_2, V)$ where Q_2 is obtained from P' by deleting the prefix of length i and the suffix of length $|P'_2| - i_2$.

Proposition 7.3 *The computation of solution sets is correct: for each derived query* $(r(i, P', j), b')$ *we have* $S_D(r(i, P', j), b') = Sol_D(r(i, P', j), b')$.

The proof is a simple modification of the earlier correctness proof. We just use Proposition 7.2 instead of Proposition 4.7.